Nocturnal visual orientation in flying insects: a benchmark for the design of vision-based sensors in Micro-Aerial Vehicles

Gert Stange (Principal investigator),
Josh van Kleef (Research scientist),
Richard Berry (Ph D student),
Robert Parker (Engineer)

Research School of Biological Sciences
Australian National University
P O Box 475
Canberra 0200
Australia
Email: gert.stange@anu.edu.au

Report Documentation Page

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
09 MAR 2011	**Final**	**20-09-2007 to 20-09-2010**

4. TITLE AND SUBTITLE	
Nocturnal visual orientation in flying insects: a benchmark for the design of vision-based sensors in Micro-Aerial Vehicles	5a. CONTRACT NUMBER **FA48690714057**
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)	
Gert Stange	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
Australian National University,PO Box 475,Canberra ACT 0200,Australia,NA,NA	**N/A**

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
AOARD, UNIT 45002, APO, AP, 96338-5002	**AOARD**
	11. SPONSOR/MONITOR'S REPORT NUMBER(S) **AOARD-074057**

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT
One of the prerequisites for achieving autonomous flight, in machines as well as organisms, is the maintenance of a stable attitude relative to the vertical direction. In order to obtain attitude information, it is possible to use visual cues, such as the mean direction of illumination, the position of the horizon or optic flow. Recently, visual attitude stabilization has been implemented in a number of designs for micro-aerial vehicles. It is also being used by flying insects, and many of them carry ocelli, a supplementary set of simple lens eyes that appear to be specialized for attitude stabilization. These ocelli are particularly well developed in dragonflies, which are among the most competent fliers in existence. We have explored the dragonfly ocellar system, identifying the spatiotemporal transfer functions of its components, namely the dioptric system, the photoreceptor neurons, and, most importantly, the secondorder neurons which are the final stages in processing visual information. The outputs of a total of 15 neurons form a map of part of the visual world, covering a wide range (200?) of azimuth along the horizon. In elevation, the range is narrow (30?) and different elevations are not mapped onto different neurons, but the outputs of all neurons contain a component that is sensitive to optic flow in elevation. Thus, it appears that the dragonfly ocellar system uses an approach that distinguishes it from existing technical horizon detection algorithms. Our knowledge about details of the system is sufficient to implement it in hardware and to commence testing in realistic situations.

15. SUBJECT TERMS
Biomimetics, Vision, Insect Flight Control, Micro Air Vehicles (MAVs)

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**44**	

Contents

Summary

One of the prerequisites for achieving autonomous flight, in machines as well as organisms, is the maintenance of a stable attitude relative to the vertical direction. In order to obtain attitude information, it is possible to use visual cues, such as the mean direction of illumination, the position of the horizon or optic flow. Recently, visual attitude stabilization has been implemented in a number of designs for micro-aerial vehicles. It is also being used by flying insects, and many of them carry ocelli, a supplementary set of simple lens eyes that appear to be specialized for attitude stabilization. These ocelli are particularly well developed in dragonflies, which are among the most competent fliers in existence. We have explored the dragonfly ocellar system, identifying the spatiotemporal transfer functions of its components, namely the dioptric system, the photoreceptor neurons, and, most importantly, the second-order neurons which are the final stages in processing visual information. The outputs of a total of 15 neurons form a map of part of the visual world, covering a wide range (200°) of azimuth along the horizon. In elevation, the range is narrow (30°) and different elevations are not mapped onto different neurons, but the outputs of all neurons contain a component that is sensitive to optic flow in elevation. Thus, it appears that the dragonfly ocellar system uses an approach that distinguishes it from existing technical horizon detection algorithms. Our knowledge about details of the system is sufficient to implement it in hardware and to commence testing in realistic situations.

Introduction

Airborne organisms or machines are, in principle, manoeuvrable in six degrees of freedom, three translational ones, defining the position, and three rotational ones, defining attitude. In order to maintain stable flight it is necessary to have some means of controlling these variables.

The present report is specifically concerned with two components of attitude, namely rotation around the lateral axis, referred to as pitch, and rotation around the longitudinal axis, known as roll. In both cases, the prerequisite for control is knowledge of a reference vector defining the vertical direction. Such knowledge cannot be derived from gravitation sensors, because such sensors cannot discriminate between acceleration forces and gravitation. For example, the rotational acceleration whilst flying a banked turn results in a false vertical signal which is displaced around the roll axis (Fig. 1), and linear forward acceleration affects the pitch reading. Thus, the airframe must either carry gyroscopes with negligible drift, which is not always feasible, or it must use an external reference such as the visual environment.

The visual environment contains a number of cues that could be relevant for attitude stabilization, for example, in ascending order of complexity,
- the 'centre of gravity' of mean overall illumination which, in an outside environment, can be close to the vertical direction,
- the position of the boundary formed by the horizon,
- the rate of image movement (optical flow) in different parts of the visual field. This requires higher level image processing and is not suitable for defining an absolute reference.

The need for vision-based horizon detection

•In flying organisms or machines, acceleration sensors that detect the direction of gravity give a false reading if additional linear or angular acceleration forces are present, for instance during a turn.
•Therefore, the horizon becomes a dominant cue for the maintenance of equilibrium. In order to maintain level flight, it is sufficient to make sure that the horizon bisects the visual field in all directions.
•Details of the scene are irrelevant, the resolution of a panoramic horizon detector can be quite low and the detector can be quite simple.
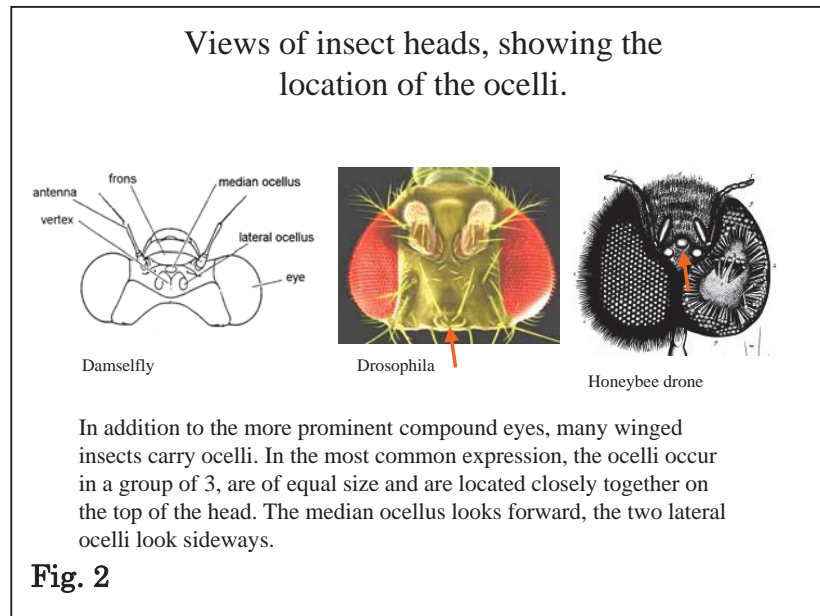
Fig. 1

The present report addresses the extent to which an insect visual system, formed by the ocelli of dragonflies, encodes these different aspects of the visual environment within the context of pitch and roll attitude stabilization.
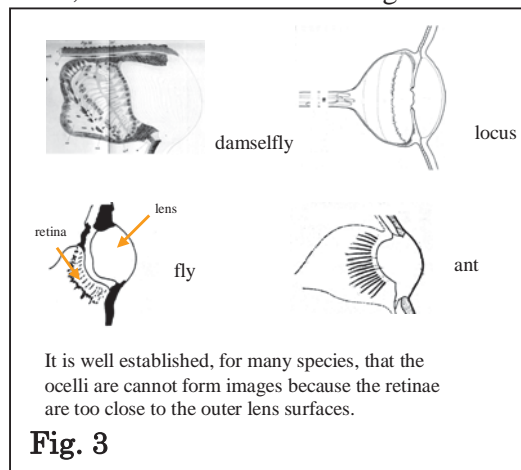
This report endeavours to provide a comprehensive coverage of the field, exceeding the scope of our specific project in some cases, but including most of our key findings, some previously reported and some new. The overall aim is to derive a blueprint for a biologically inspired technical attitude controller; however, it is only natural that there has been a two-way flow of information, leading to the incorporation of technology-inspired concepts into our understanding of the biological system. Consequently, we include an overview of developments in the design of vision-based attitude controllers, a field that has been undergoing rapid growth in the recent past, driven by widespread interest in control systems for micro-aerial vehicles. An application to micro-aerial vehicles would also be the next stage of development that may follow on from this report.

Previous work on structure and function of insect ocelli

In addition to the more prominent compound eyes, many winged insects carry ocelli. In the most common expression, the ocelli occur in a group of 3, are of equal size and are located closely together on the top of the head. The median ocellus looks forward, the two lateral ocelli look sideways (Fig. 2). As early as 1908, this latter observation led the German zoologist Rudolf Hesse to the suggestion that those ocelli are involved in horizon detection.

Views of insect heads, showing the location of the ocelli.

Damselfly

Drosophila

Honeybee drone

In addition to the more prominent compound eyes, many winged insects carry ocelli. In the most common expression, the ocelli occur in a group of 3, are of equal size and are located closely together on the top of the head. The median ocellus looks forward, the two lateral ocelli look sideways.
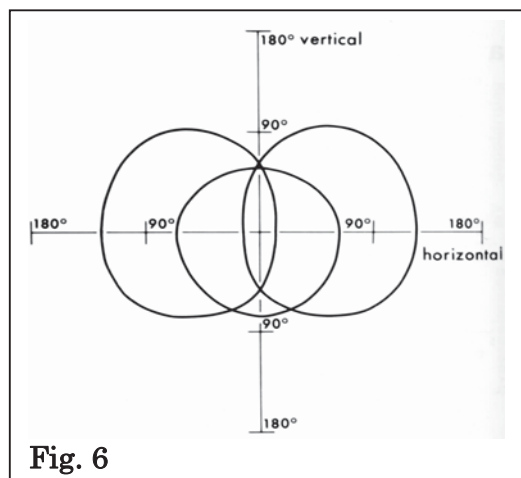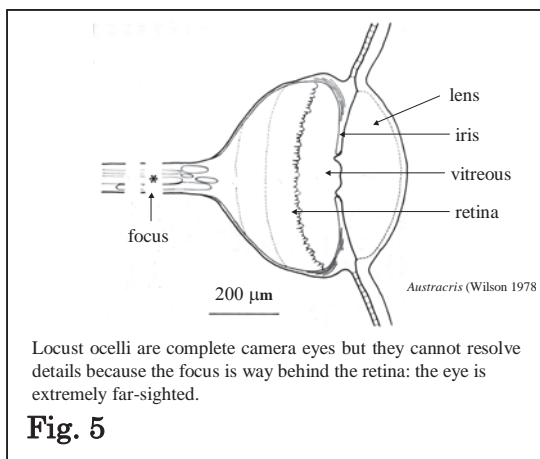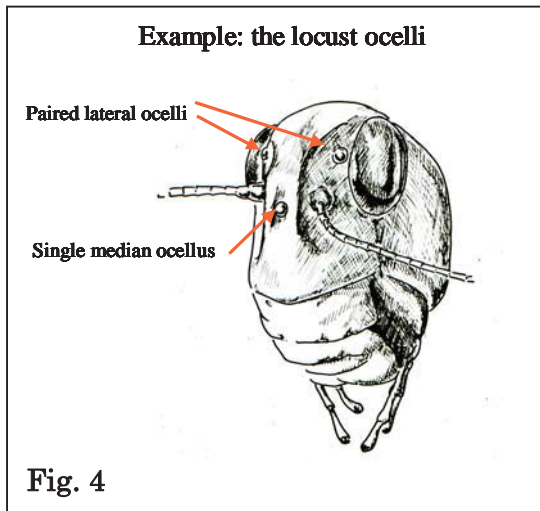
Fig. 2

In the following decades, a substantial body of work on ocelli was created by a considerable number of authors, on anatomy as well as physiology and behaviour. There were many suggestions as to what the function of the ocelli might be, but no general consensus was reached, mainly because there was no clear-cut behavioural response: when the compound eyes are painted over, insects behaved as if they are blind, but the effects of blinding the ocelli were manifold and ambiguous. Given the large number of insect species, possessing ocelli but living in different habitats, it is reasonable to assume anyway that there is no single ocellar function.

damselfly

locus

lens

retina

fly

ant

It is well established, for many species, that the ocelli are cannot form images because the retinae are too close to the outer lens surfaces.

Fig. 3

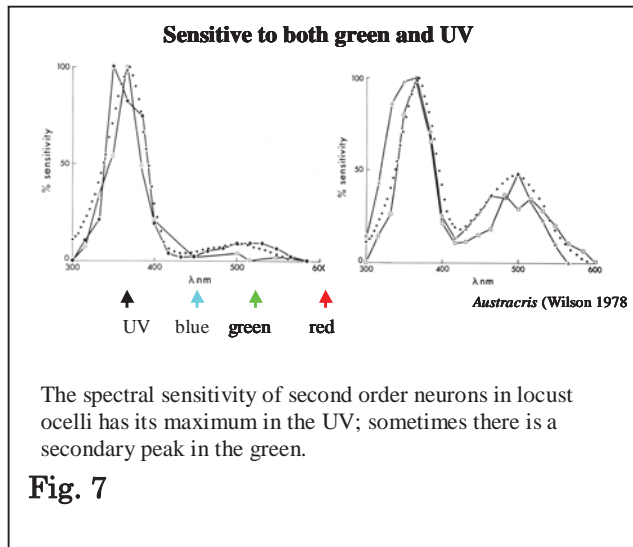By circa 1970, a partial consensus was reached, recognizing that ocelli are particularly strongly expressed in insects that can fly, and that the optics of ocelli are such that image formation is impossible (Fig. 3): Although the lenses are strongly curved, the refractive index of the lens material is too low (approx. 1.5) for image formation at the position of the retinae. Instead, images are positioned way behind the retinae.

Example: the locust ocelli

Paired lateral ocelli

Single median ocellus

Fig. 4

lens

iris

vitreous

retina

focus

200 μm

Austracris (Wilson 1978)

Locust ocelli are complete camera eyes but they cannot resolve details because the focus is way behind the retina: the eye is extremely far-sighted.

Fig. 5

180° vertical

90°

180° | 90° | 90° | 180°

horizontal

90°

180°

Fig. 6

The horizon hypothesis. The first hint that Hesse's conjecture was right came from observations in locusts (Goodman 1965), showing that obscuring the ocelli reduces the accuracy with which those insects maintain a flight attitude in which they turn their backs towards a light source (dorsal light response). Fig. 4 gives a portrait of a locust head, showing the positions of the compound eyes and the ocelli. Wilson (1978) verified that locust ocelli are also underfocused (Fig. 5); he supplemented this observation by electrophysiological recordings from the large second-order neurons (L-neurons) in the ocellar retinae, measuring the electrical signals evoked by light flashes from different directions (angular sensitivity functions) and by light flashes of different wavelengths (spectral sensitivity functions).

The angular sensitivity functions are extremely wide but distinct for different ocelli, consistent with the finding that the optics are underfocused (Fig. 6). Taken together the fields of view (FOV) of the three ocelli cover nearly the full viewsphere. During level flight, half of each FOV is covered by ground and the other half by sky.

Sensitive to both green and UV

Austracris (Wilson 1978)

UV   blue   **green**   **red**

The spectral sensitivity of second order neurons in locust ocelli has its maximum in the UV; sometimes there is a secondary peak in the green.
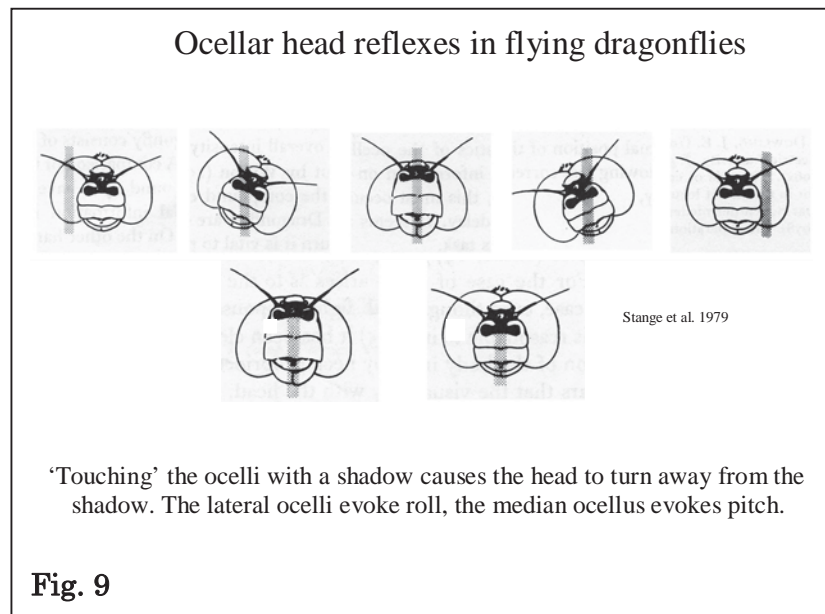
Fig. 7

The spectral sensitivity of L- neurons in locust ocelli (Fig. 7) has its maximum in the UV; sometimes there is a secondary peak in the green. When we look at a scene in the UV (Fig 8), it becomes evident that using UV for horizon detection is an excellent idea because most objects and surfaces on the ground are quite dark at this wavelength range, whereas the sky is blue, radiating predominantly at short wavelengths including UV. In the green the sky-ground contrast is particularly low, but a green channel could be used as a reference to obtain invariance against changes in light intensities at different times of day.

Images of a scene at two different wavelengths

UV   Green

Wilson 1978

Using UV for horizon detection is an excellent idea because the sky-ground contrast is high. In the green it is particularly low, but a green channel could be used for normalization.

Fig. 8
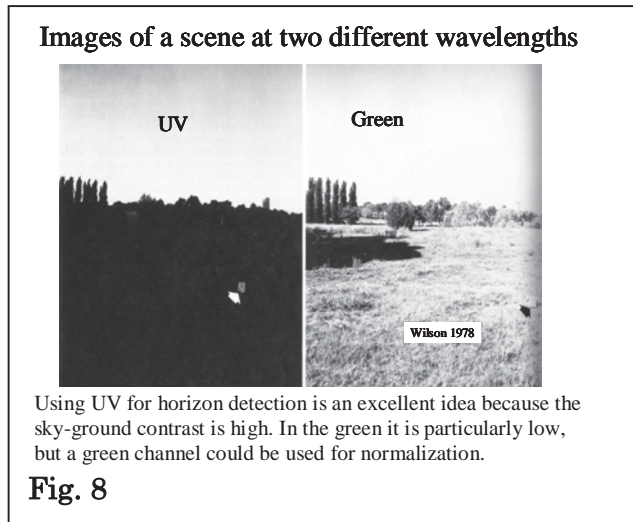
Behavioural verification. Wilson's conclusions encouraged a study by Taylor (1980), who went to the extreme step of surgically removing the compound eyes in early instars of locusts with the hope that the animals would survive and moult into adults. He succeeded and was able to demonstrate the presence of dorsal light responses in the absence of compound eyes. Around the same time, Stange (1979) used a less invasive method on dragonflies. The experimental animal is attached to a rod by its thorax, with the head freely movable. It is illuminated from a point source such that all ocelli receive some light. When exposed to an air flow, the animal will go into flight mode. Under those conditions, head reflexes can be evoked by the appropriate changes in illumination, and those reflexes are evoked by the ocelli. When the shade of a matchstick is cast over either ocellus (Fig. 9), we observe corrective head movements, with the apex of the head turning away from the shaded ocellus, rotating around the roll axis on shading either lateral ocellus (first row in Fig. 9) and rotating around the pitch axis on shading the median ocellus (second row). Shading the compound eye has no effect.

With the independent findings by Taylor and Stange the question of ocellar function seemed to be settled. Few papers addressed the topic in the following 25 years. Amongst those are papers by Rowell (1988) who demonstrated the presence of

ocellar inputs to flight motorneurons in locusts, and a paper by Schuppe and Hengstenberg (1990) who demonstrated that ocellar head reflexes are present in flies, albeit of small amplitude.

Ocellar head reflexes in flying dragonflies

Stange et al. 1979

'Touching' the ocelli with a shadow causes the head to turn away from the shadow. The lateral ocelli evoke roll, the median ocellus evokes pitch.

Fig. 9

Technical horizon sensors

Over the past few years, a remarkable proliferation of designs for micro-aerial vehicles (MAVs) has occurred, driven by a combination of technological progress and strong interest by defence agencies (Fig. 10). This has, in turn, created interest in
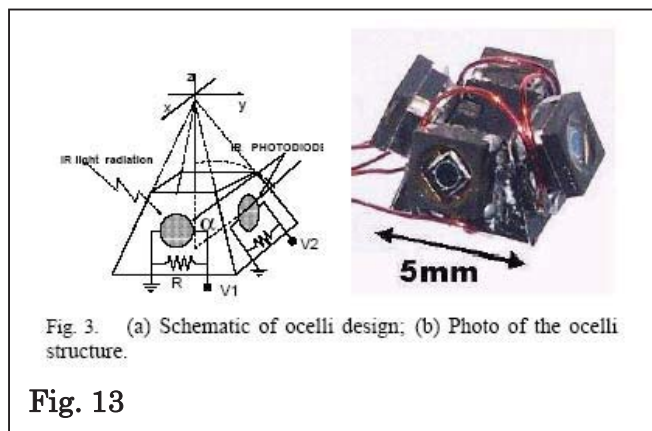

There is an impressive diversity of designs, indicative of the strong interest in the field.
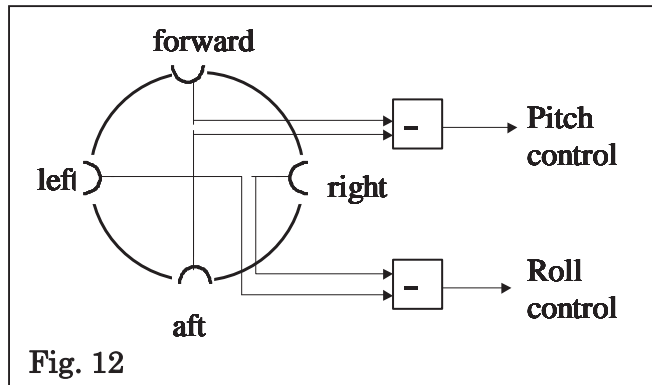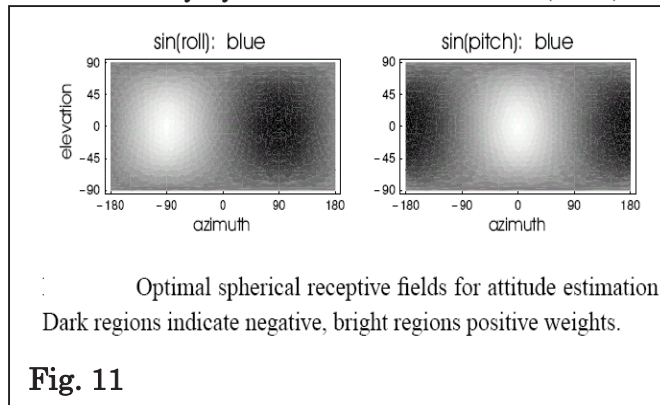
Fig. 10

lightweight attitude control systems, some of them vision-based. Some designs are biologically inspired, others have emerged independently. The surge of interest in vision-based horizon sensors is attributable to the realization, now widespread, that MAV flight stability presents particularly difficult challenges. To quote from Ettinger et al. (2001):

'The first step in achieving such MAV autonomy is basic stability and control, although this goal presents some difficult challenges. The low moments of inertia of MAVs make them vulnerable to rapid angular accelerations, a problem further complicated by the fact that aerodynamic damping of angular rates decreases with a reduction in wingspan. Another potential source of instability for MAVs is the relative magnitudes of wind gusts, which are much higher at the MAV scale than for larger aircraft. In fact, wind gusts can typically be equal to or greater than the forward airspeed of the MAV itself. Thus, an average wind gust can immediately affect a dramatic change in the flight path of these vehicles.'

Two approaches are being used for the implementation of technical horizon detectors for MAVs. In one approach, the overall centre of incident illumination (mean luminance vector) is determined by two wide-field optical sensors, one each for pitch and roll, essentially reducing the task to simple analog hardware. This is very fast but limited in accuracy. In the other approach, the horizon line is extracted, by machine vision algorithms, from a stream of video images taken from the MAV. The task involves two-dimensional, multi-pixel image processing, making it accurate but computationally intensive.

Detecting the mean luminance vector. This approach uses purely reactive mechanisms based on optimized receptive fields of small numbers of detector elements, incorporating prior knowledge only. For the case of attitude control, it was shown formally by Neumann and Buelthoff (2002) that a pair of wide field detectors, assigning two orthogonal sinusoidal weight functions to the full viewsphere (Fig.11), is optimal for the detection of roll and pitch. As acknowledged by the authors, the spatial sensitivity distributions resemble those of insect ocelli, as described above. In Fig. 12, we show the block diagram of a simple hardware implementation. Four wide field sensors are aligned with the horizontal plane. The idea is that a tilt leads to intensity differences between opposite sides because the ground is darker than the sky. Simply feeding the difference signals from each pair to the control surfaces of the aircraft will close feedback loops for roll and pitch. Note that three detectors would be enough, albeit computationally more complex.

A fly-ocelli inspired attitude controller (Fig. 13) was designed by Schenato et al. (2004), and a simple, photoresistor-based autopilot for model aircraft has been commercially available for so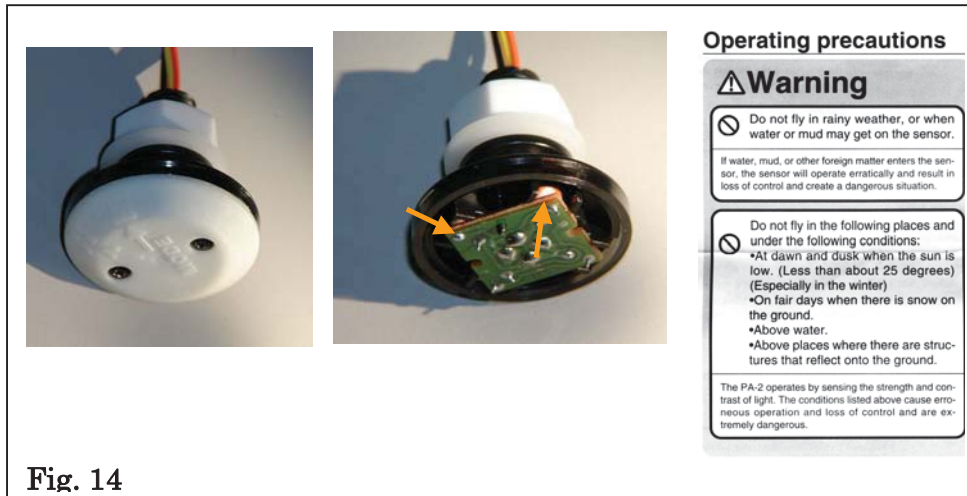me time (Futaba PA-2, (Fig. 14)). The latter is of limited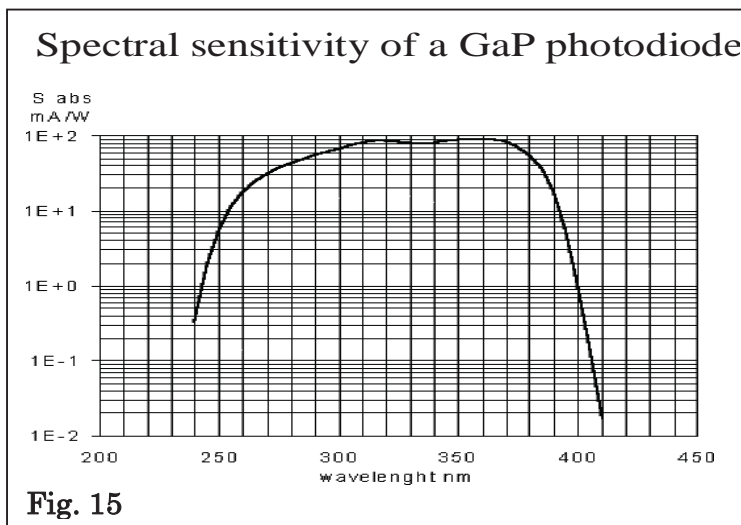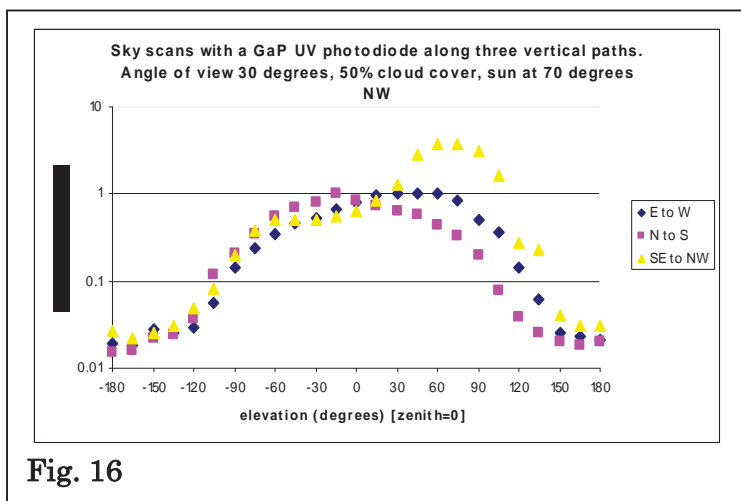 usefulness, as it is sensitive to the presence of bright surfaces on the ground and to a low sun. It is also limited in the range of usable light intensities. The main reason for this is probably its use of green-sensitive detectors, the wavelength at which sky-ground contrast is lowest.

Optimal spherical receptive fields for attitude estimation Dark regions indicate negative, bright regions positive weights.

Fig. 11

Fig. 12

Fig. 3. (a) Schematic of ocelli design; (b) Photo of the ocelli structure.

Fig. 13

Fig. 14

Operating precautions

⚠ **Warning**

🚫 Do not fly in rainy weather, or when water or mud may get on the sensor.

If water, mud, or other foreign matter enters the sensor, the sensor will operate erratically and result in loss of control and create a dangerous situation.

🚫 Do not fly in the following places and under the following conditions:
• At dawn and dusk when the sun is low. (Less than about 25 degrees) (Especially in the winter)
• On fair days when there is snow on the ground.
• Above water.
• Above places where there are structures that reflect onto the ground.

The PA-2 operates by sensing the strength and contrast of light. The conditions listed above cause erroneous operation and loss of control and are extremely dangerous.

Fig. 15
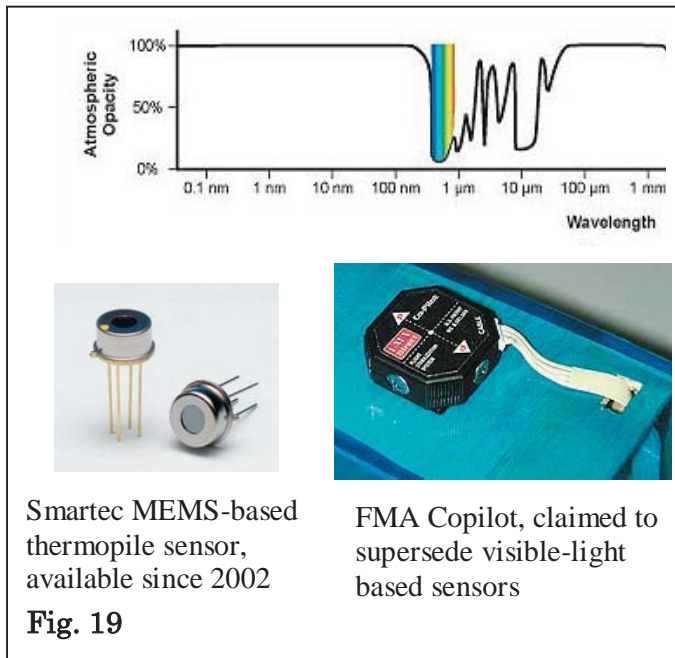
Spectral sensitivity of a GaP photodiode

Fig. 16

We tested Wilson's conjecture (Fig. 8) that the usage of UV is advantageous by using a Gallium phosphide photodiode (available since 1999), with a spectral sensitivity that resembles the spectral sensitivity of a locust ocellus (Fig. 15), and an angular FOV of ±30° in two dimensions. The detector was mounted on a sky scanner and

photocurrents were plotted along 3 scans through 360°, along an east-west trajectory, a north-south trajectory, as well as a trajectory from southeast to northwest for part of which the full sun would directly shine into the detector (Fig.16). It is evident that the response functions are roughly sinusoidal, but that exposure to the sun (at 70° in yellow line) causes a major irregularity, although it did not drive the sensor into saturation. Notice that a logarithmic scale is used, consistent with the fact that the intensity-response functions of biological photosensors often show a compressive nonlinearity.

Fig. 17

Pairs of green and UV photodiodes

Fig. 18
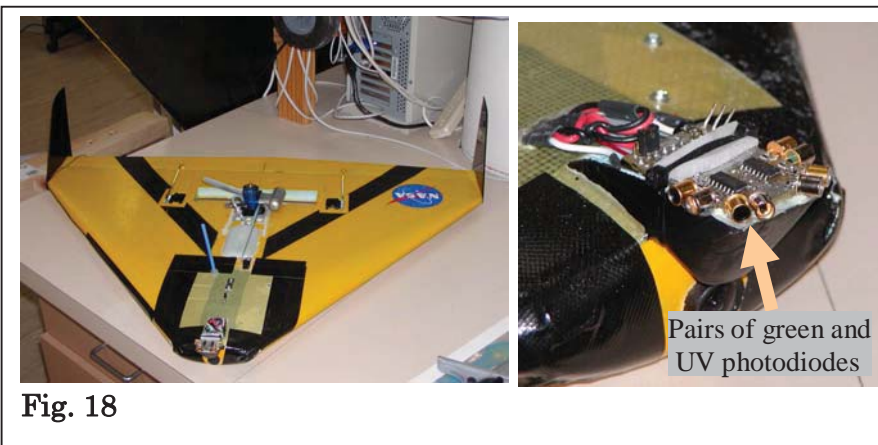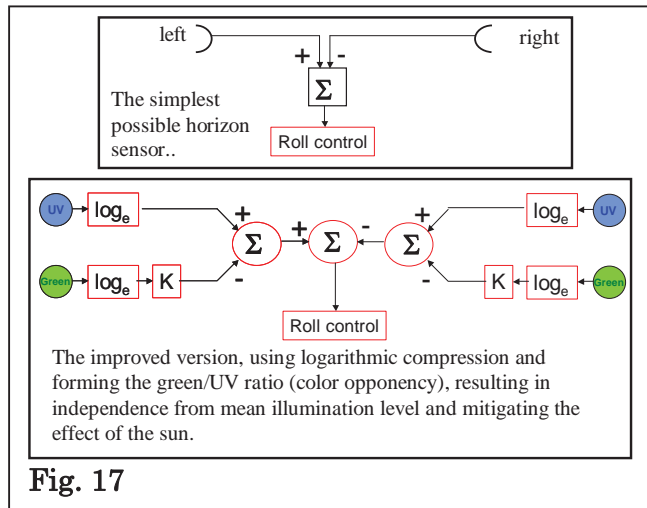
Inspired by ocellar vision, we optimized the optical-wavelength approach (Chahl et al. 2003), by using pairs of UV and green sensitive detectors, by using logarithmic compression and by forming the ratios of the output pairs (colour opponency) (Fig. 17). Thereby, the system was supposed to become robust against fluctuations of overall light intensity and the effect of a low sun was supposed to be mitigated. This attitude controller was incorporated into a MAV that was also supplied to AFRL, Eglin Air Base (Fig.18); its success, if any, is not known to the present authors.

Smartec MEMS-based thermopile sensor, available since 2002

FMA Copilot, claimed to supersede visible-light based sensors

Fig. 19

Optical-wavelength based attitude controllers can be extremely small and have response times in the order of microseconds, but they suffer from the disadvantage that they depend on ambient light from an external illuminant. The atmosphere is opaque to most wavelengths below 1 mm, except for windows in the visible and infrared ranges (Fig. 19). In the 10 µm IR window, the ground is always brighter than the sky because it is warmer, independent of the presence of external illumination. Therefore, a device using this window will work at night and will be insensitive to the colour of the ground and to a low sun.

Seen at MAV'05 conference, Garmisch 2005

Fig. 20

Recently, mass-produced thermopile detectors have become available, with response times in the order of tens of milliseconds. As a result, their usage has found its way into the commercial market, and FMA-Direct (Fig. 19) is marketing an IR-

based attitude stabilizer for model aircraft, claiming performance that is far superior to that of optical-wavelength systems. At the 1st US-European MAV competition, September 2005, the MAV group at the University of Toulouse, France, gave flight demonstrations of MAVs equipped with an IR-based horizon sensor (Fig. 20). They worked well, but for the fixed wing aircraft quite noticeable oscillations aro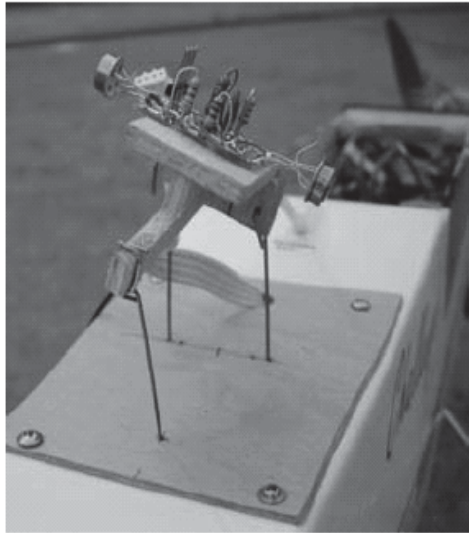und the roll axis were observable, at a frequency of a few Hz. This could be due to the slow response characteristics of the IR sensors, something that is likely to become more important with decreasing size. This could be an intrinsic shortcoming of the IR approach, meaning that it would be premature to dismiss the UV approach on which the present paper is focused.

The merits of the IR-based approach for the control of MAVs are also being investigated by, for example, Taylor et al. (2003)(Fig. 21), who mounted a pair of roll-detecting sensors on a platform which is kept in line with the horizon by a servo. Thus, the error signal to this servo becomes the control signal for roll stabilization.


Fig. 21

Detecting the horizon boundary. Human pilots detect the horizon by recognizing one part of the visual world as sky and another as ground. In formal terms, this process can be described as classifying all points on the visual sphere into either sky or ground, by criteria such as texture and colour, and then estimating the best separating line between those two classes. This classical machine vision approach incorporates multi-element imaging photodetector arrays, as well as computationally demanding stages such as pixel classification and image segmentation. It may also incorporate an adaptive learning step whereby the rules for picture element classifications are determined immediately prior to usage.

Machine vision approach to horizon detection
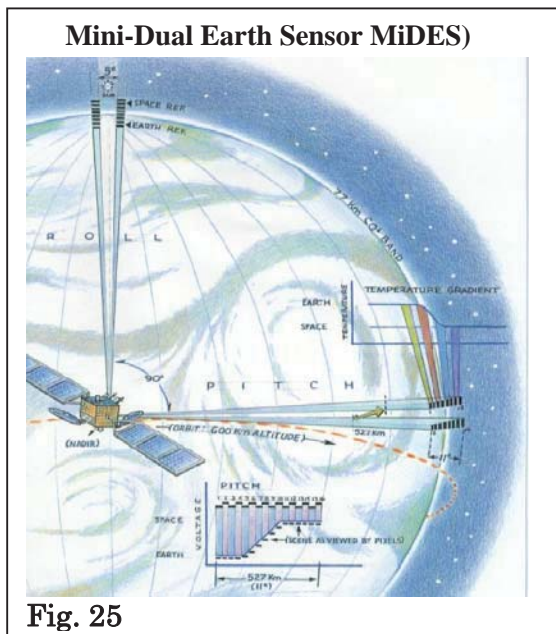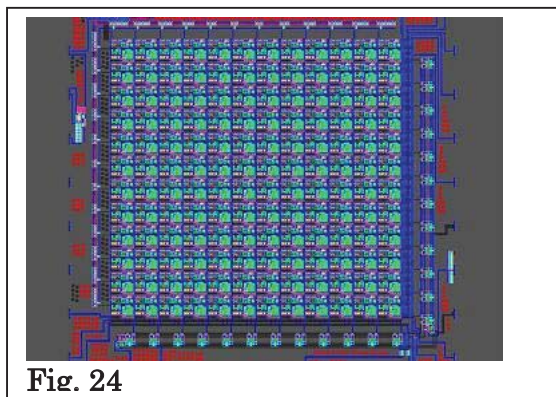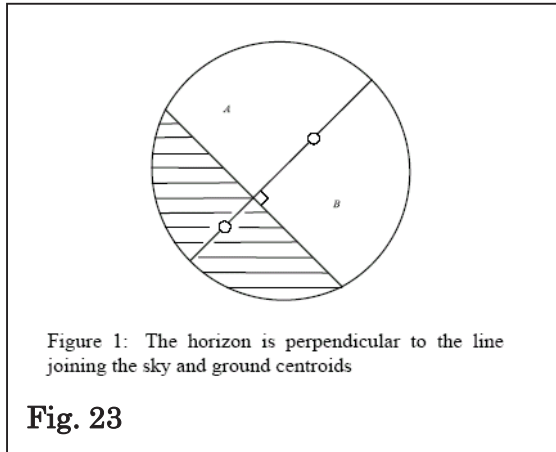(after Ettinger et al., Univ. Florida (2004))

₂Any line through the image divides it into two classes of pixels (above and below the line).
₂By assuming that colours of sky pixels will be closer to those from other sky pixels than to those from the ground (and vice-versa), then the best estimate of the horizon is given by the line that produces minimum overall variation in each of the two groups.
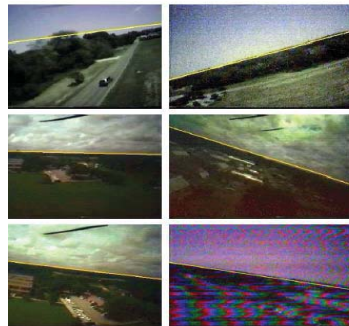
Various horizon-detection examples under different lighting conditions (sunny and cloudy), and with varying degrees of video transmission noise. Yellow lines: best estimates.

Fig. 22

Figure 1: The horizon is perpendicular to the line joining the sky and ground centroids

Fig. 23

Fig. 24

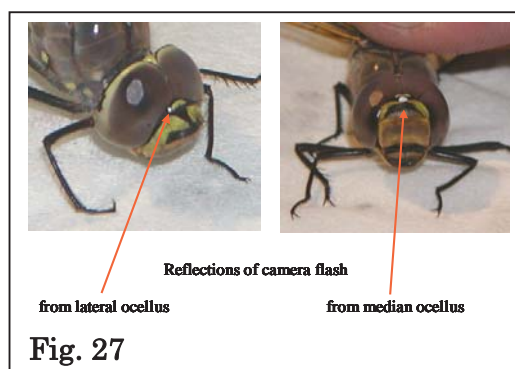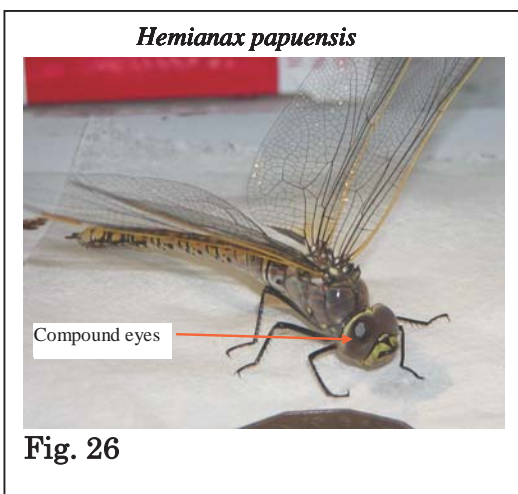**Mini-Dual Earth Sensor MiDES)**

Fig. 25

Such algorithms were implemented, by processing a video stream from a camera on a serial computer, by Ettinger et al. (2004) (Fig 22) and Winkler et al. (2004). Cornall et al.(2005) present an interesting idea on how to simplify computations (Fig. 23): In a circular image, slope and position of a straight horizon boundary can be computed by simply calculating the centroids of the two classes of pixels. Horiuchi (4) designed a CMOS based, analog VLSI 12*12 photoreceptor array that finds a best-fit horizon line by parallel processing (Fig. 24).

A standard method of attitude stabilization for orbiting satellites (Fig. 25) uses one-dimensional arrays of IR-sensors and is an interesting variant of the boundary detection approach.

To summarize this chapter, it appears that technical horizon detectors use either two pairs of very wide field detectors, with a panoramic view, or an array of multiple detectors which look only at a small fraction of the total viewsphere. The first method is computationally simple but prone to luminance irregularities such as the sun, whereas the second method is computationally demanding and accurate for a distant horizon at large altitude, but it must fail when large objects are within the field of view, such as in an urban jungle. This gives rise to the suggestion that it might be possible to merge both approaches. As will be demonstrated in the following chapters, it appears that the dragonfly ocellar system is doing exactly this.

16

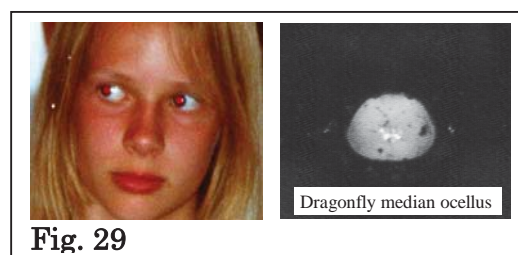# The dragonfly ocelli: fields of view and optics

Figure 26 gives an overall view of one of the dragonfly species we used, *Hemianax papuensis*, with a wingspan of about 80 mm. In close-up views (Fig. 27),

Hemianax papuensis

Compound eyes

Fig. 26

Reflections of camera flash

from lateral ocellus   from median ocellus

Fig. 27

Compound eyes   Vertex

Lateral ocellus   Lateral ocellus

Median ocellus

Fig. 28

Dragonfly median ocellus

Fig. 29

the ocelli light up from the flash of the camera, a lateral one when viewed from the side and the median one when viewed from the front. As seen in a close-up view (Fig. 28), the ocellar system in dragonflies is unusual in that the median ocellar lens is oblong and quite large (0.6 mm across in *Hemianax*) and that it is recessed below a protruding vertex. The lateral ocelli are positioned on the sides of the vertex.
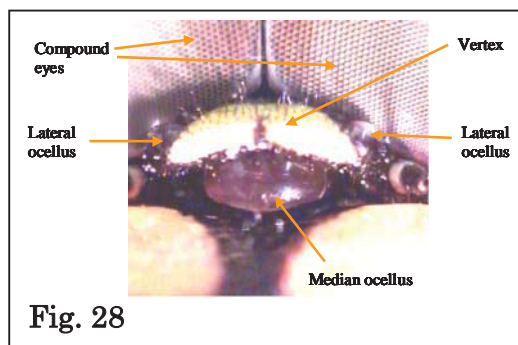
The eyeshine phenomenon as observed in Fig. 27 is attributable to reflections from the inside of the eye, similar to the 'red eye' phenomenon which is a nuisance in portrait photography (Fig. 29); the ocellar eyeshine is quite bright because the ocelli contain a reflecting tapetum like found in cats' eyes. We can use this phenomenon as a non-invasive method to determine the fields of view of the ocelli: if light from a point source is reflected from the inside of an eye, the source must be visible to the subject in the first place. We can also use the phenomenon for testing whether light is focused on the tapetum (and therefore the retina): if this is the case, only a small patch will reflect back, and because of the reversibility of optics it will be visible only over a small range of angles.

The ophthalmoscope shown in Fig. 30 allows inspection of the ocellus from all directions and is designed to meet both purposes. A small LED light source can either be attached to the camera, axially in front of the lens, or it can be in a fixed position. If the light source is moved with the camera, we can obtain the overall fields of view, regardless of whether the eye is focused or not. If the light source is fixed, eyeshine from a focused eye will only be visible if the direction of view is close to the direction

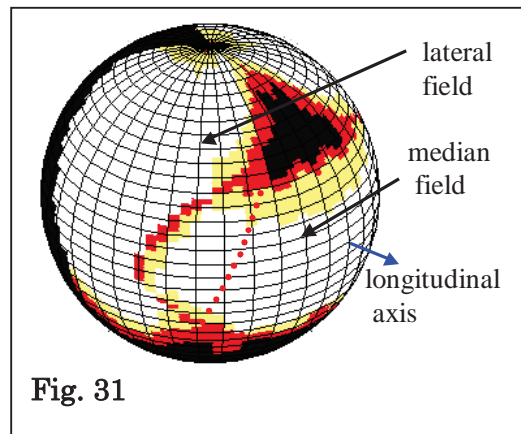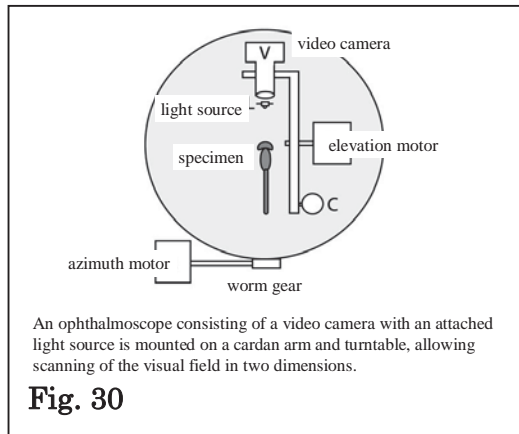of illumination but it will be visible over a wide range if the eye is not focused.

An ophthalmoscope consisting of a video camera with an attached light source is mounted on a cardan arm and turntable, allowing scanning of the visual field in two dimensions.

Fig. 30

Fig. 31

Figure 31 shows the overall ocellar fields of view for the dragonfly *Hemicordulia tau*, projected on a sphere. Directions from which eyeshine is clearly visible are marked in yellow or white, red indicates the boundaries and black marks the direction from which no eyeshine was visible from either ocellus. The lateral and median fields overlap partially and the median field is drawn on top of the lateral one. The forward boundary of the latter is indicated by a dotted red line.

We find that the median ocellus covers a forward looking field that is centered on the longitudinal axis; its dimensions are anisotropic, extending over ±60° in azimuth and ±30° in elevation. The lateral fields look sideways and cover 30° to 130° in azimuth; in elevation, they extend from the zenith to -40°, way below the equator.

Figure 32 shows the result of the same experiment performed in the locust. It is evident that the ocellar receptive fields are much wider and there is a much wider overlap. The finding is consistent with the observation by Wilson shown in Fig. 6.


Fig. 32

Fig. 33
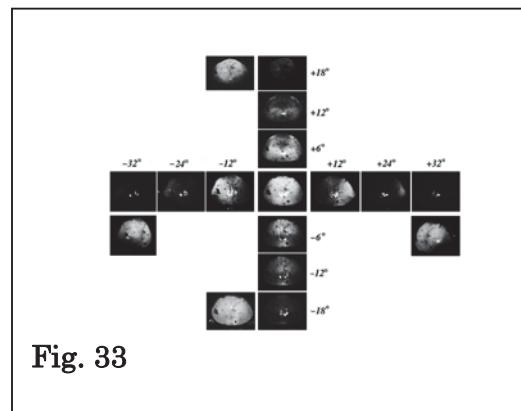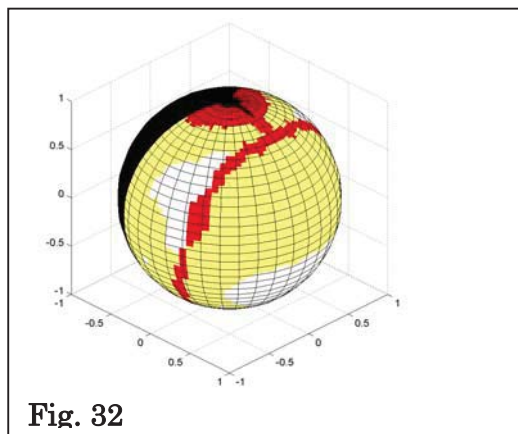
In Figure 33 the light source was in a fixed position, aligned with the longitudinal axis of the head of a dragonfly and images were taken from different directions, to determine the angular range for which a fixed eyeshine patch is visible. We find that the eyeshine patch vanishes once we go off-axis by more than 6° in


Fig. 34

Fig. 35
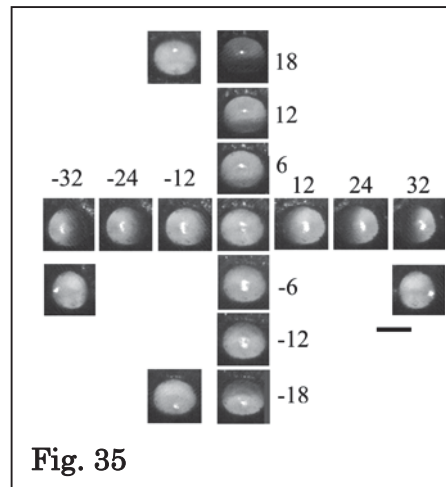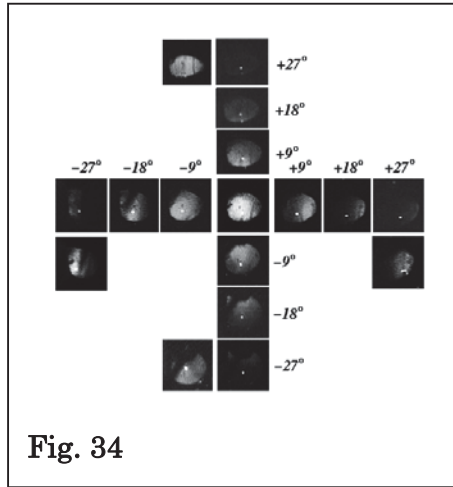
elevation or 12° in azimuth. This demonstrates that the dragonfly median ocellus forms a focused image on the tapetum, It also indicates some anisotropy, with a resolution in elevation that is double that in azimuth. The extra figures next to the arms of the cross are controls where the light source was axial to the camera.

In Figure 34, the result of the same experiment on a lateral ocellus is shown, again demonstrating the presence of a focused image. Under the same conditions, evidence for focusing is absent in the locust (Fig. 35), again consistent with the observations by Wilson (Fig. 6).

Fig. 36

Fig. 37
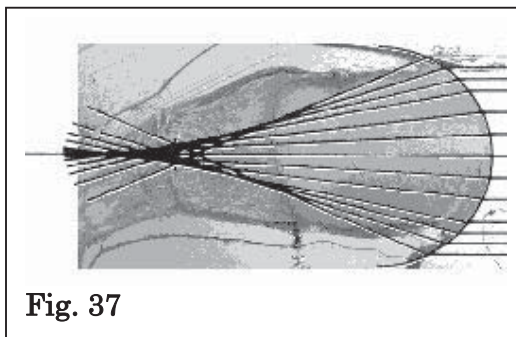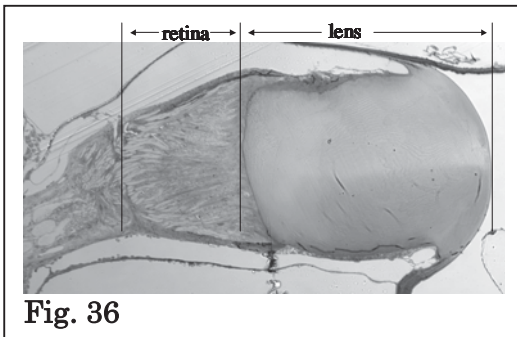
Further information on the optical properties of the dragonfly ocelli was obtained using microscopy. A vertical cross-section through the median ocellus (Fig. 36) reveals that the lens is unusually thick, when compared to other ocelli (Fig. 3). In Figure 37, rays are traced through the frontal surface of the lens, assuming a homogenous refractive index of 1.5. Under this assumption, the focus would be just beyond the retina with a considerable depth of uncertainty due to spherical aberration. However, interference microscopy (which shows different refractive indices as different colours) reveals the presence of a refractive index gradient (Fig. 38), a property that compensates for spherical aberration and results in a shorter focal length.

Fig. 38

sagittal   transverse

Fig. 39

Fig. 40
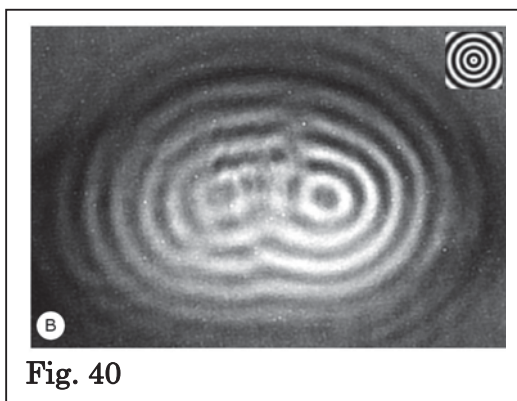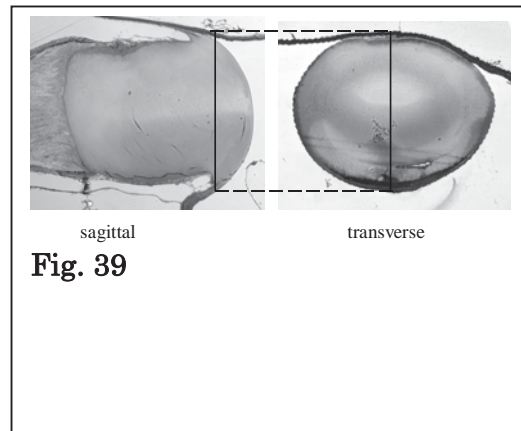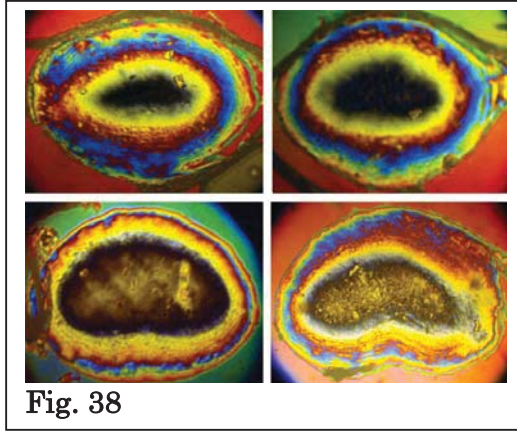
As reported above, eyeshine measurements indicate different vertical and horizontal resolutions, and this is confirmed by anatomical studies. The shape of the lens is anisotropic (Fig. 39), leading to the expectation that vertical and horizontal contours are imaged in different ways. Viewing the image that the lens forms of a set of concentric circles confirms this (Fig. 40). We also used a more precise method, namely determining the position of best spatial cutoff frequency by looking at images of horizontal and vertical gratings. It is found that the focal planes are within the retina (Fig. 41).

The spatial cutoff frequency of the lens is 1.2 cycles/degree in the vertical.

Posterior limit of retina

Analysis of images of horizontal and vertical gratings shows that the image planes formed through the astigmatic lens are positioned within the retina, 100 and 180 μm respectively, behind the back surface of the lens.

Fig. 41
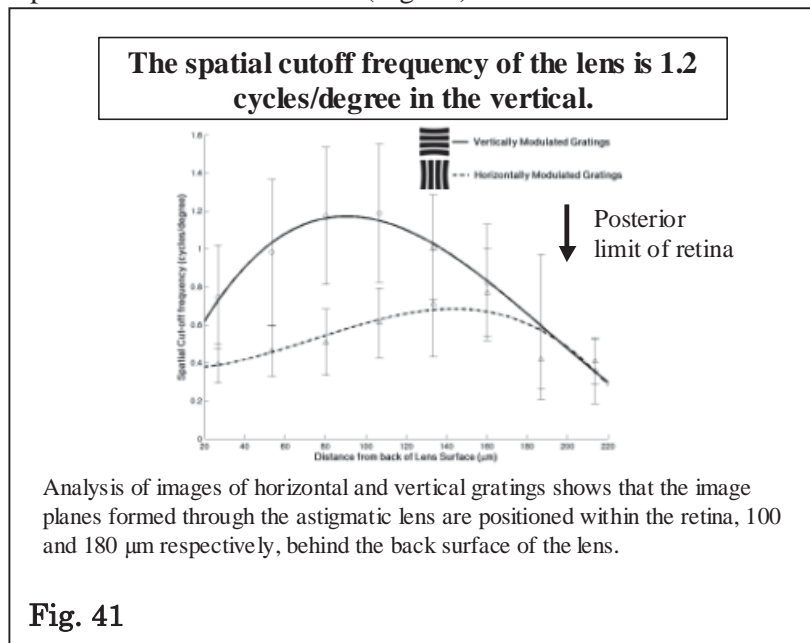
The anatomy of the lateral ocelli (Figs. 42-44) turns out to be more complex. Basically, the lens is bifocal and there are two separate retinae, dorsal and ventral. The image of the lower part of the FOV, projected on the dorsal retina, is in focus, the

image of the upper part is not. Thus, we expect that the dorsal retina can resolve details, like the dragonfly median ocellus, while the ventral retina sees a blurred image, as is the case in the ocelli of other insects.

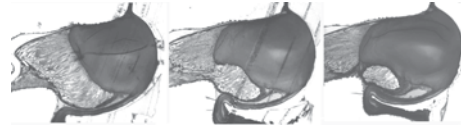**Hemicordulia tau, approximately horizontal sections of right lateral ocellus**
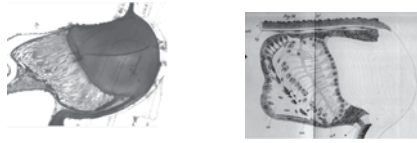
ventral

Agrion (Hesse 1902)

In this section, the geometry of the dragonfly lateral ocellus is quite similar to what has long been known for damselflies: the focal plane is beyond the retina….

Fig. 42

Hemicordulia tau, approximately horizontal sections of right lateral ocellus

ventral   median   dorsal

….but this is not the case for sections at other levels. Proceeding from ventral to dorsal we first encounter a step. In the dorsal part, the focal plane is on the retina. There are two distinct retinae and neuropiles.

Fig. 43

Hemicordulia tau, 3-D reconstruction, by confocal microscopy, of the right lateral ocellar lens

Dorso-posterior surface
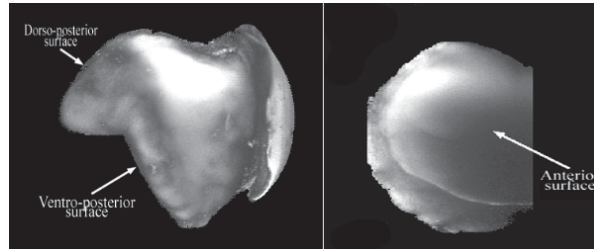
Ventro-posterior surface

Anterior surface

Fig. 44

The dragonfly ocelli: transfer functions of the receptor neurons

The optics of the ocelli project a partially processed image of the world onto the receptor neuron layer, and the transfer characteristics of that layer define the next information processing stage of the system. In particular, we are interested in how the ocellar response changes with mean intensity, temperature, angular position and wavelength. Furthermore, we want to examine the extent to which the system is linear, and the extent to which individual elements interact with their neighbours.

We decided to adopt the computationally demanding methodology of presenting pseudorandom UV and green spatiotemporal patterns followed by Wiener kernel analysis, for several reasons. First, the method is suitable to measure responses to modification of several stimulus dimensions concurrently, and it allows us to detect and determine nonlinear interactions. Second, earlier efforts to determine spatial transfer functions by sequential measurement of responses to light flashes from a point source provided no clear results because the receptor neurons, and more so the second order neurons, contain an automatic gain controlling mechanism (adaptation). The independent, random modulation of a spatial array of light sources is necessary to avoid this problem.

Fig. 45

A one-dimensional, multi-element visual display was used for the measurement of the photoreceptor characteristics. Sixteen pairs of green and UV LEDs are positioned on the circumference of a circle (Fig. 45), covering 75°, such that their optical axes intersected at the centre of the circle. The arrangement is attached to a cardan arm, allowing manual rotation perpendicular to the optical axes of the LEDs. This allows the display to be used at multiple angles as either a vertical or a horizontal one-dimensional display.

Each LED is independently driven via a 32-channel D/A converter, which in turn is driven by a computer program that could transfer data at a maximal rate of 3 megabytes/s. The PC is run on the Linux operating system, with a real-time module (RTIA) for synchronous data acquisition and control, interfaced with Matlab® for higher level functions. For electrophysiological recordings, the LEDs were driven by a 1-dimensional pseudorandom pattern of 20 s duration, at a refresh rate of 625 Hz. The physiological response of the neuron was recorded in synchrony with the stimulus.

Time course, spectral sensitivity, temperature dependence. The time course of the responses (Fig. 46), in response to a light impulse, consists of a latency (typically 8 ms at 23° C), presumed to be attributable to the phototransduction process, followed by the sum of a proportional (first order) term and a differential (second order) term, presumed to be attributable to the dynamics of ion channels in the receptor neuron

22

membrane and to intracellular biochemical processes. As biological processes are based on chemical reactions, they tend to be strongly temperature dependent. For the function of sensory systems, this is of little consequence in warm-blooded animals, but in insects such as dragonflies the ability to regulate temperature is quite limited. Yet, the temperatures at which dragonflies can fly range from below 20° C to above 40°C.
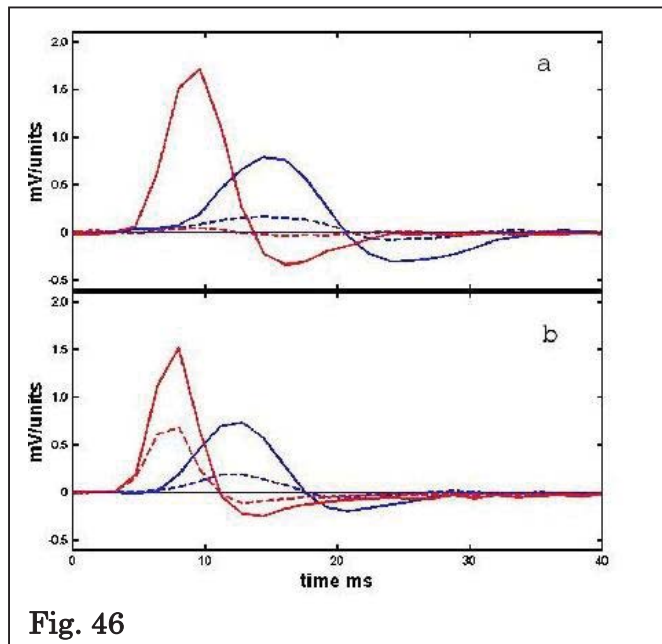

Fig. 46

Dragonfly ocellar photoreceptor neurons contain a UV and a green photopigment and it is also known from the literature that different receptor neurons vary in their relative UV and green sensitivities. Hence, we routinely determined UV/green ratios. At a temperature of 23° C, we found that UV/green sensitivity ratios varied between 1.2 and 21 (n = 37, mean = 6.9, SD = 6.3).

We measured temperature dependencies of parameters of ocellar receptor neurons, finding no significant effect on angular sensitivities, but there is a pronounced effect on time course and spectral sensitivities. In both examples of Fig. 46, latencies and times to peak decrease markedly when the temperature is raised from 23° C (blue) to 31° C (red), while the response amplitudes increase. Unexpectedly, there is also a strong but heterogeneous effect on relative UV and green sensitivities. In example (a) the response to green (dotted lines), at 23 ° C, is approximately 25% of the UV response (solid lines) but it becomes barely detectable at 31° C. In example (b), the response to green is larger at 31° C, both in relative and absolute terms.

Angular sensitivity: elevation.
Fig. 47 shows a first order kernel (equivalent of the impulse response) of a receptor neuron, for the dimensions of elevation and time. The recording was obtained after manually positioning the one-dimensional array at the azimuth of maximum sensitivity to brief flashes of light from all LEDs in the array. We notice that this particular neuron had its maximum response at an elevation of 20° above the horizon; on either side of this maximum, the response falls off rapidly. In the time domain, there is a delay (latency) of 8 ms before onset of a positive-going (depolarizing) response, which reaches its maximum after 14 ms. The depolarizing component is followed by hyperpolarization.

The same procedure was repeated for a total of 16 receptor neurons. Their

angular sensitivities (defined as the response widths at 50% of maximum) and locations in azimuth were mapped on the outline of the total field of view of the
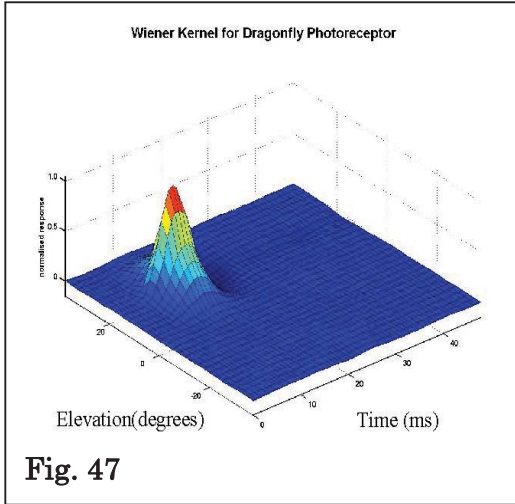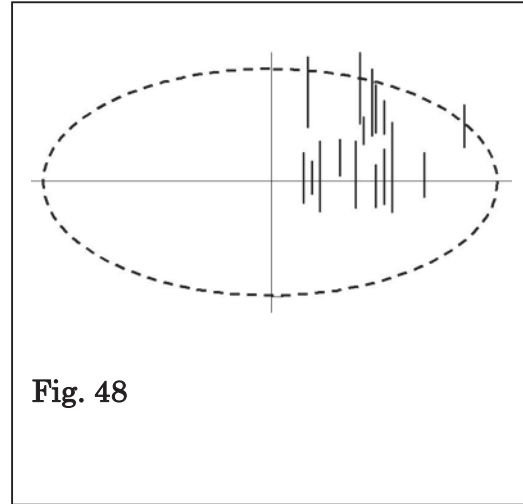

Fig. 47

Fig. 48

ocellus (an ellipse,120° wide and 60° high), as shown in Fig. 48. The mean angular sensitivity is 14.7°, with a standard deviation of 4.3°. All responses occurred to the left of the vertical midline, consistent with the fact that all recordings were made from the right retina. There is a bias towards the upper half of the receptive field, which we attribute to bias in the recording procedure: The electrode was always inserted into the retina from below, meaning that upwards-looking receptor neurons were encountered first.

Angular sensitivity: azimuth.
The data represented in Figs. 49 and 50 were derived in a similar way as those in the preceding section, except that the array was mounted horizontally, on a cardan
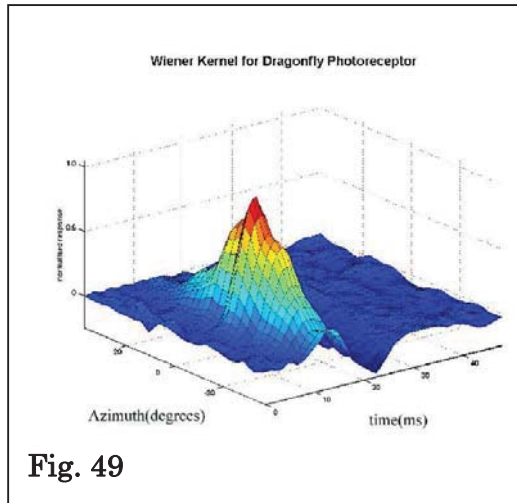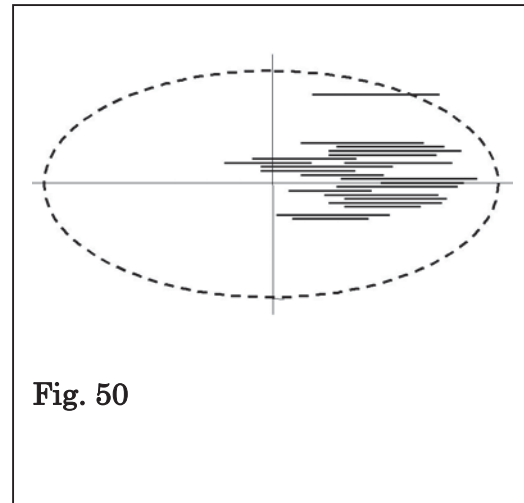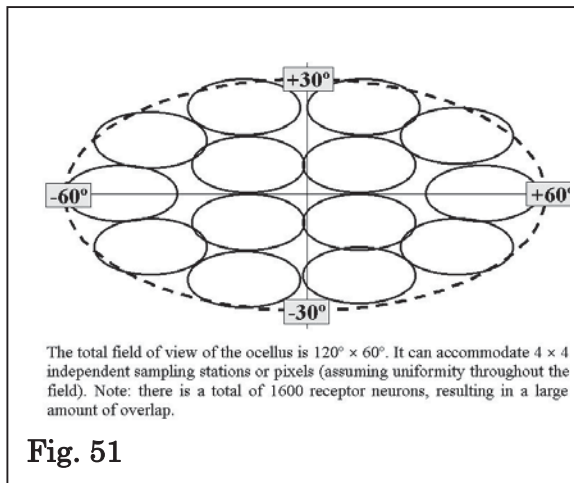

Fig. 49

Fig. 50

arm that allowed manual adjustment to the maximum response in elevation, prior to application of the stimulus. Analogous considerations apply; mean angular sensitivity was 27.6°, with a standard deviation of 4.6°.

The total field of view of the ocellus is 120° × 60°. It can accommodate 4 × 4 independent sampling stations or pixels (assuming uniformity throughout the field). Note: there is a total of 1600 receptor neurons, resulting in a large amount of overlap.

Fig. 51

In the eyeshine measurements, the spatial resolution of the median ocellus was found to be approximately 12° in elevation and 24° in azimuth (Fig. 33). The current findings, by the most direct method possible, confirm this estimate. They also confirm the earlier conclusion that the fields are elliptical, with an axis ratio of 2. This leads to the inference (Fig. 51) that the overall field of view can accommodate 4 pixels in elevation and 4 pixels in azimuth, or a total of 16 independent sampling stations. However, that inference needs to be qualified: The total number of receptor neurons is approximately 1500, implying spatial oversampling by a factor of 100. Such oversampling could, for instance, improve signal-to-noise ratio at low intensities or be relevant for fast motion detection.

The dragonfly ocelli: anatomy of the second-order neurons

Given that the outputs of the receptor layer encode a 2-dimensional image of the environment, the question immediately follows as to how this information is mapped onto the following layer, the second-order neurons (L-neurons). As the

Fig. 52
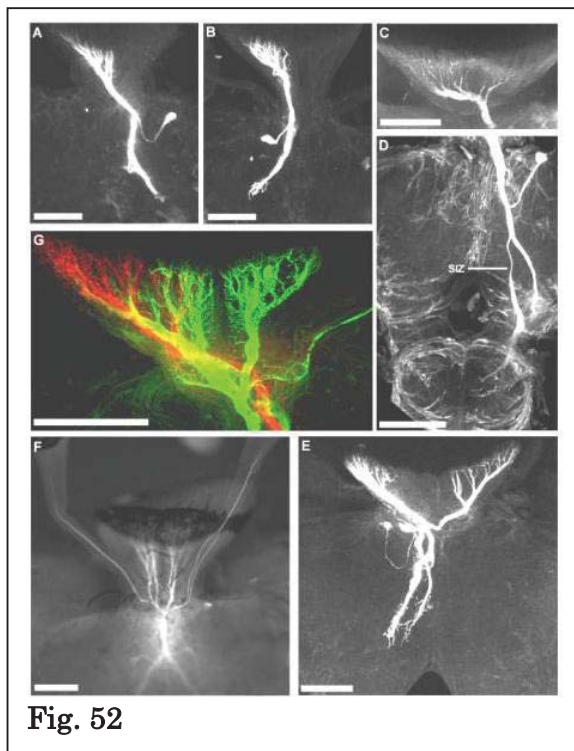
number of independent channels at the receptor neuron outputs is 16 and the number of second-order neurons (as known prior to the present study) is roughly the same, namely 9, a naïve expectation would be that there is simply a congruent mapping, whereby each L-neuron samples a separate patch of retina, such that different L-neurons all look in different directions, both in elevation and azimuth. One way to examine this hypothesis is to investigate the anatomical branching patterns of the L-neurons in the retina.
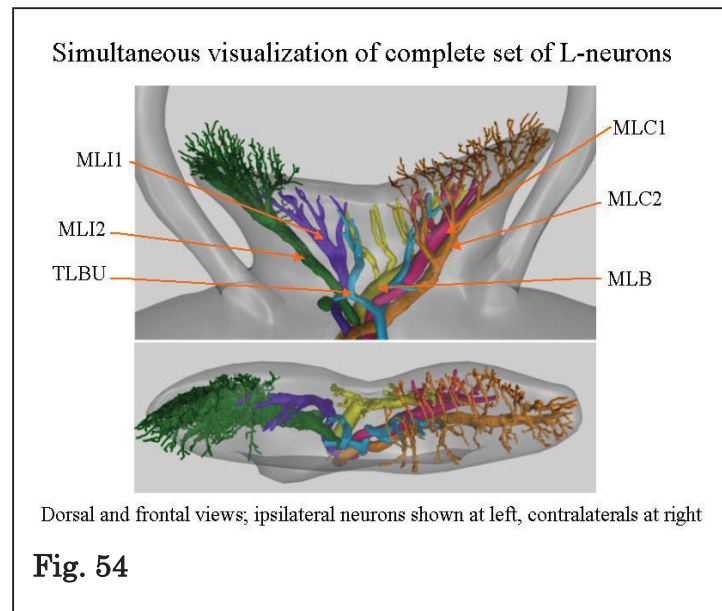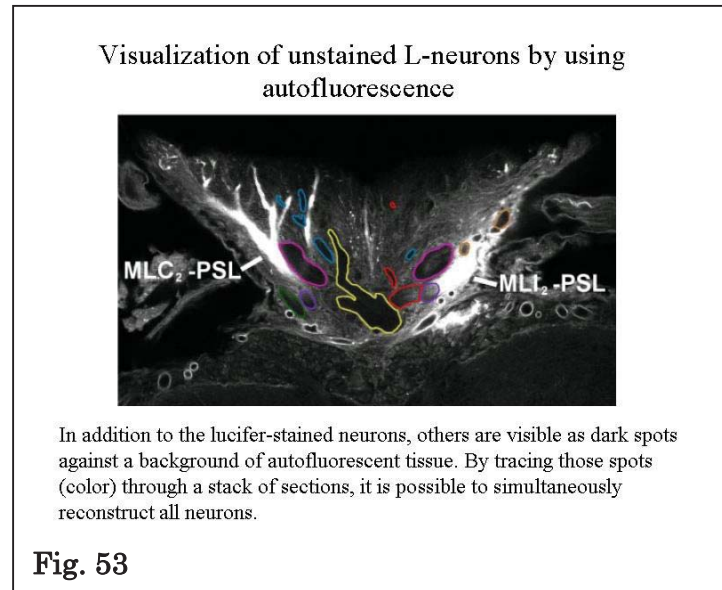
By applying a modern staining method (Lucifer Yellow iontophoresis), we have been able to determine the branching patterns for the complete set of L-neurons, in three-dimensional reconstruction, using two independent methods,

namely confocal microscopy and reconstruction from serial sections.

Using confocal microscopy, Fig. 52 shows ventral views of stained median ocellar L-neurons in the dragonfly brain from six different preparations of *Hemicordulia tau* or *Hemicordulia australiae*. All types of L-neuron present in the median ocellar nerve are represented and were identified anatomically by iontophoretic injection of dye after recording from the unit. (A) MLC_1-PSL, one of a pair of neurons which originate from one half of the ocellar neuropile and project to the contralateral posterior slope (PSL). (B) MLI_1-PSL, similar to MLC_1-PSL except that the axon of this neuron terminates in the PSL on the same side of the brain as the dendritic branching. (C) The dendritic branching pattern of one of the paired MLB-D neurons. These neurons branch to both halves of the ocellar neuropile but in an asymmetric fashion, such that one half of the ocellar neuropile is more heavily innervated than the other. The axon projects to the half of the brain contralateral to the side of the ocellar plexus which contains the most prolific branching. As seen in (D) the axon of MLB-D has at least two sites of termination: one in the posterior deutocerebrum and one projecting at least as far as the suboesophageal ganglion (SOG). The presumable spike initiation zone is shown (SIZ). (E) A double fill of two types of paired neurons. MLC_2-PSL is shown on the right side of the ocellar neuropile and MLI_2-PSL on the left. The axon of MLC_2-PSL projects to the contralateral half of the brain with respect to the dendritic branches, while the axon of MLI_2-PSL remains within the ipsilateral half. The dendritic branching of MLI_2-PSL is limited to the lateral extremes of the ocellar plexus. (F) The unpaired, bilaterally symmetrical TLBU-PSL neuron. Branches of this neuron project to all three ocelli and innervate the median ocellar plexus and PSL on both sides of the brain. (G) shows the dendritic branching of TLBU-PSL in the median ocellus at high resolution after staining with Lucifer Yellow (green). A second cell MLC_1-PSL has also been filled with Rhodamine Dextran (red) showing the two neurons branching to different regions within the ocellar neuropile. Note that a third cell (MLC_2-PSL) has been partially filled with Lucifer Yellow which gives the impression of asymmetry across the midline of the brain in the green channel. With the exception of (F) all images are maximum intensity projections from a series of optical sections taken through the brain. (F) is an image taken under direct epifluorescence. Scale bars for all figures are 200 μm.

The alternative method is tracing through a stack of sections, using the fact that the axons of individual neurons, if unstained, do not autofluoresce, whereas the surrounding tissue does. Reconstructions were performed on a single preparation of *Hemicordulia tau* from which two neurons had been stained by iontophoretic injection of Lucifer Yellow; the remainder was traced using different colours. Fig. 53 shows a ventral view of an optical section through the median ocellar nerve. The two stained neurons, MLC_1-PSL on the left and MLI_2-PSL on the right, are clearly visible as strongly fluorescent structures. Autofluorescent footprints of unstained L-neurons are shown as coloured contours. To avoid ambiguity the left and right members of the bilateral MLB-D neurons are shown as yellow and red respectively. Other paired neurons innervate one half of the median ocellus only and therefore both members of the pair are shown in the same colour. Reconstructions of the stained and unstained neurons within the median ocellus are shown from a ventral viewpoint in Fig. 54 (top) and a frontal viewpoint in Fig. 54 (bottom). For clarity only the left member of each

pair of neurons is shown (where left refers to the location of the axon or cell body of the neuron within the brain from a dorsal view).

Visualization of unstained L-neurons by using autofluorescence

In addition to the lucifer-stained neurons, others are visible as dark spots against a background of autofluorescent tissue. By tracing those spots (color) through a stack of sections, it is possible to simultaneously reconstruct all neurons.

Fig. 53

Simultaneous visualization of complete set of L-neurons

Dorsal and frontal views; ipsilateral neurons shown at left, contralaterals at right

Fig. 54

In summary, we observe that the second-order neurons have distinct individual characteristics, both in the branching patterns in the retina and the target areas in the brain. We also observe that there is a strong overlap in the horizontal direction (mapping azimuth), although none of the second-order neurons have branching patterns that cover the whole horizontal extent of the retina. However, regarding the vertical projection (mapping elevation), there is no evidence that different neurons project to different regions. Rather, it appears that the full vertical extent of the retina is sampled by all neurons.

The dragonfly ocelli: spatial transfer functions of the second-order neurons

The core component of the present project consists of the determination, by the usage of electrophysiological methods, of the way the second order neurons encode visual information in space and time. In the preceding paragraphs we have resolved three prerequisites, namely the determination of the transfer characteristics of the optics and of the receptor neuron layer, as well as the morphological identification and classification of the second order neurons. The latter results enable us to electrophysiologically characterize one neuron at a time and then determine its identity by electrophoretic staining.

Stimulus: two-dimensional array.

As we expected that the properties of the second-order neurons would be much more complex than those of the receptor neurons, we considered a one-dimensional stimulus display as inadequate and decided to construct a two-dimensional display. After it had become clear that we had to construct the display from discrete components, with considerable costs per pixel, it was necessary to use a design that is as parsimonious as possible. The spatial and temporal resolution of the receptor neurons obtained with the one-dimensional system gave us the necessary information to estimate the refresh rates and pixel densities required. Given that the frequency response of the receptor neurons had rolled off to 10% at 200 Hz, a refresh rate of 625 Hz was deemed to suffice. The mean half-widths of the receptive fields were found to be 14° in azimuth and 27° in elevation, and the overall size required is known from the eyeshine measurements. Consequently, a display with pixels 6° and 12° apart, in the azimuth and elevation respectively, was deemed to be sufficient.

Spatial relationship between display and ocellar field of view

The dragonfly is mounted such that most of the median ocellar field of view is covered by the display.

Fig. 55

Fig. 56
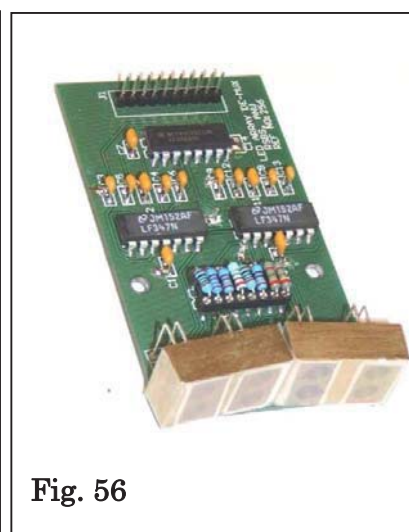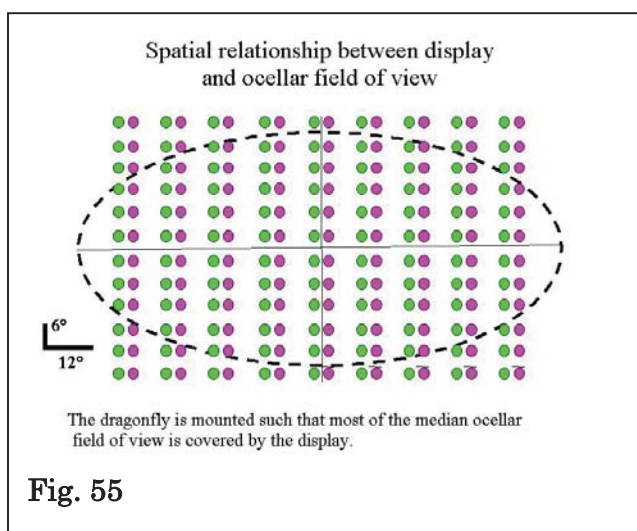
The modular design consists of 108 pairs of green and UV LEDs, mounted on 9 motherboards, each of which is equivalent to the one-dimensional array shown in Fig. 45. The overall layout, and its relation to the FOV of the median ocellus, is shown in Fig.55. Each motherboard carries 3 submodules which contain 8 LEDs (4 pairs of UV and green, Fig. 56). Each motherboard and submodule is constructed so that LEDs are separated vertically by 6° and the motherboards are inserted into a cage constructed so

they are 12° apart horizontally (Figs. 57 and 58). Thus, all individual LEDs are located at 6° intervals in the vertical and 12° intervals in the horizontal, on the surface of a sphere, covering 66° in azimuth and 108° in elevation, sufficient to sample the full field of view of the median ocellus.
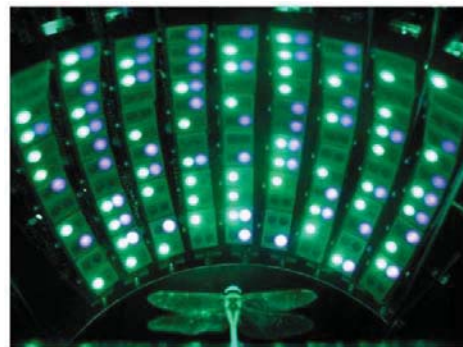

Fig. 57

Fig. 58

The driving electronics and software are straightforward extensions of those used for the one-dimensional case. The outputs of 18 D/A converters are connected, in pairs, to the 9 motherboards where they fan out vertically. They continuously and sequentially update 24 analog S/H amplifiers (8 on each of the 3 submodules).

The availability of the 2-D display enables us to determine two-dimensional receptive field maps, using the white noise approach as previously used for the determination of spatial properties of receptor neurons (van Kleef et al., 2005), after expanding it to two dimensions.

Temporal transfer functions. The temporal impulse responses shown in Fig. 59
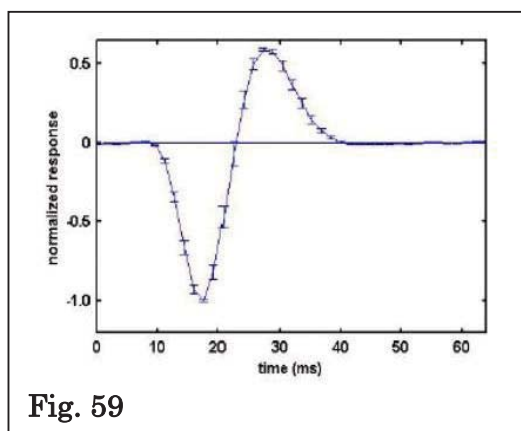
Fig. 59

is representative for the full ensemble of L-neurons. The response, measured at 22 °C, shows a delay of 10 ms, followed by a biphasic response that is characteristic for a differentiator or high pass filter. However, there is also a sustained component, as the integral is nonzero. Thus, the output of the ocellar system, as a function of intensity, contains both a proportional and a derivative component. Consequently, the system is capable of signalling the absolute position of the horizon, with an enhancement of fast changes, likely to increase stability in a closed-loop situation.
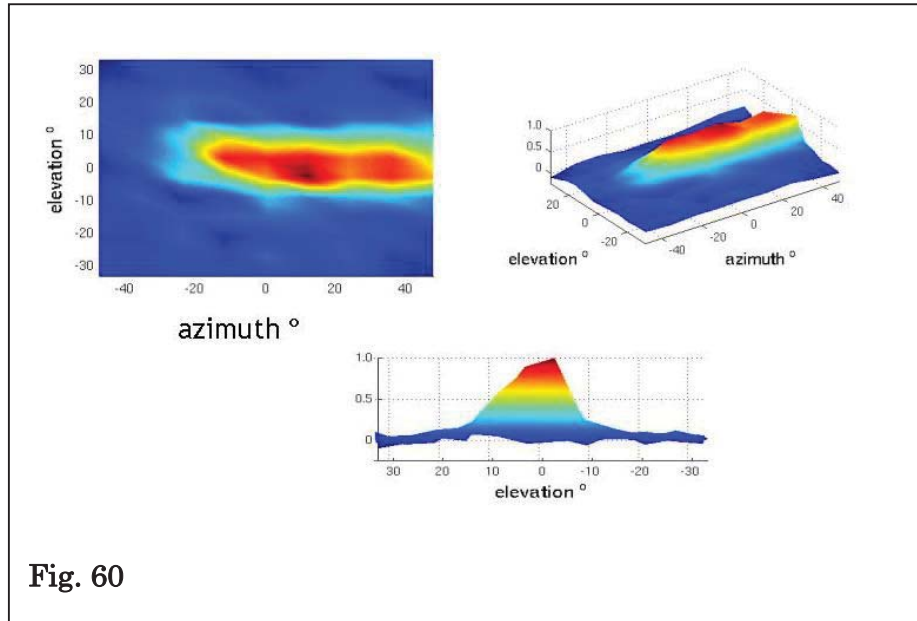
Fig. 60

Receptive fields at temporal response peak. Fig. 60 shows three different aspects of x,y,z plots of the spatial impulse response of an L-neuron, anatomically identified as MLB, stimulated by UV, at t =18 ms. The field is anisotropic. In elevation, the full width at half maximum response is 15º, centered close to the equator. In azimuth, the field extends from 25º on one side towards the full width of the ocellar field of view (60º) on the other side (clipped because our display extended to 48º only). The extent in azimuth is consistent with the anatomical dendritic branching pattern (Fig. 52 C, Fig. 54).
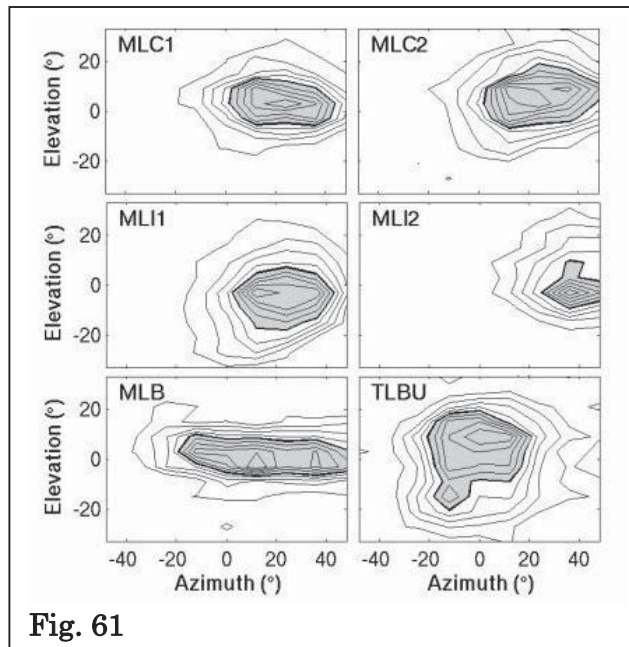

Fig. 61

In Fig. 61, contour plots of the receptive fields of all L-neuron types (for paired neurons, left members only) are shown. Fig. 62 is a schematic representation, where the lengths and widths of the smaller ellipses are mean values from several measurements on each. The large ellipse shows the extent of the total ocellar field of view, as determined from eyeshine measurements.

Fig. 62
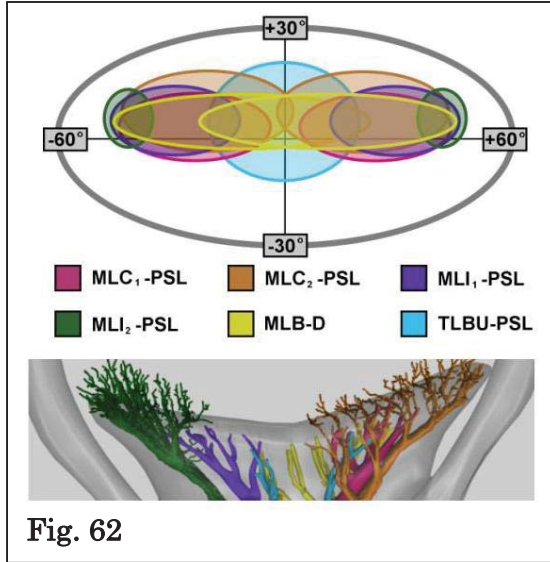
Essential features are:

- All receptive fields are close to the equator.
- In azimuth, the widths and positions of the receptive fields differ, consistent with anatomical observation. The total range is 120º.
- This geometry allows measurement of both pitch and roll.
- Every point on the equator is sampled by at least three neurons.
- In elevation, the height of individual receptive fields can be less than 20º, which is close to the acceptance angle of the receptors (15º).
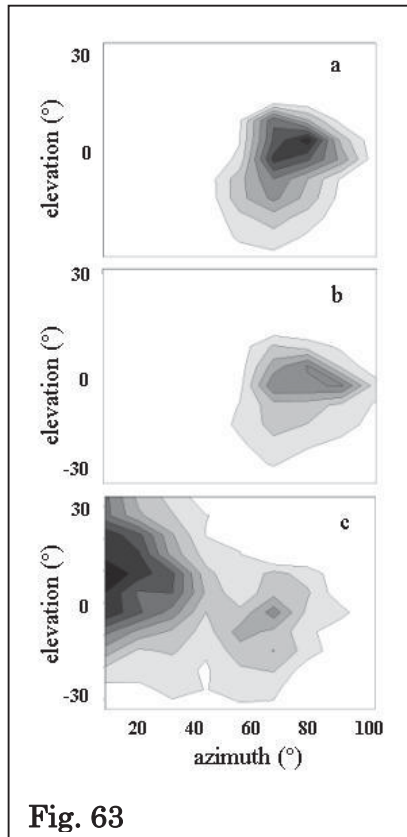
The lateral ocelli. Figure 63 shows receptive fields of L-neurons recorded in the dorsal retina of the right lateral ocellus. To obtain those measurements, the dragonfly was rotated around the yaw axis, relative to the display, by 60º. The properties of neurons (**a**) and (**b**) are consistent, in essential points, with those seen in the median ocellus: the small sizes of the receptive fields demonstrate that the dorsal retina of the lateral ocellus forms a focused image, confirming our observations in conjunction with Figs. 42-44. The neuron (**c**) is quite likely to be TLBU (Fig. 61), showing the projection of the lateral ocellar branch (Fig. 52 F) when viewed from this direction. In elevation, the directions of view are close to the equator. In azimuth the directions of view are close to 90º, confirming the expectation, from optical studies, that the lateral ocelli extend the total field of view past the 60º covered by the median ocellus. Thus, the combined sets of median and lateral L-neurons form a half-ring, covering an azimuthal range of at least 180º.

Fig. 63

The dragonfly ocelli: movement sensitivity

The description of the spatiotemporal transfer functions by separate spatial and temporal components is a first approximation only, because the spatial transfer function shows recognizable changes with time (spatiotemporal inseparability). Fig.


Fig. 64
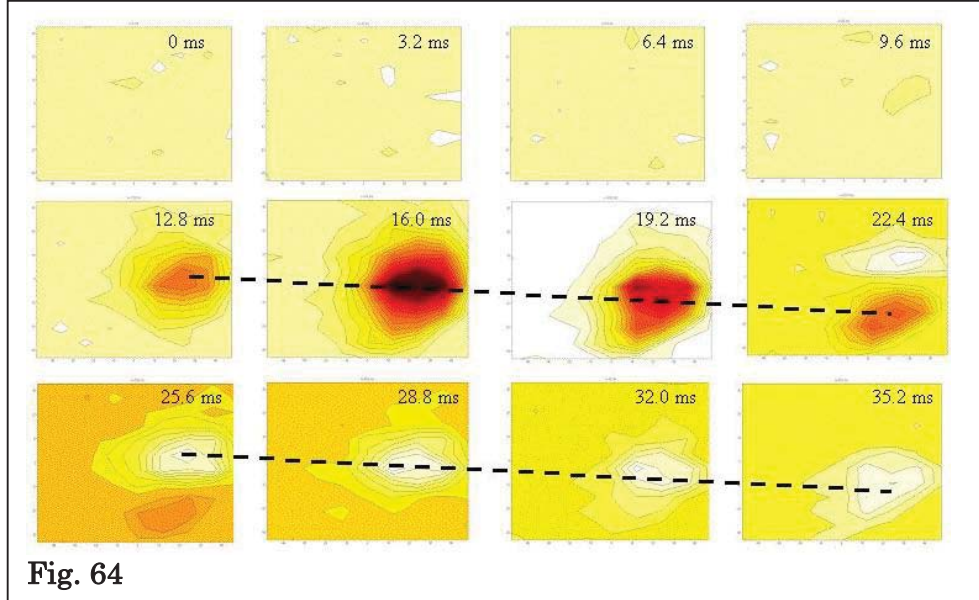
64 shows a temporal sequence of snapshots, of the spatiotemporal impulse response of one of the L-neurons, plotted in the same way as Fig. 61. The response becomes apparent at 12.8 ms and has initially a negative sign (red or black); at t = 16 ms, the largest response occurs at an elevation of -3° and an azimuth of 24°. In subsequent frames, the location of the response maximum drifts downward and the response at the original coordinates reverses its sign (white or yellow). The slope of the dotted line is the angular velocity to which the system will maximally respond. Its value is approximately 1000° / s.

Therefore, we can predict that this neuron is directionally selective to vertical movements of a horizontal edge. For an edge that traverses the field from below, at the correct speed, the responses from the two maximum positions will add, but for movement in the opposite direction the responses partially cancel each other. The polarity of the response depends on the polarity of the edge, whereby a light-dark transition evokes depolarization (the same as darkening the overall field), whereas a dark-light transition evokes hyperpolarization.

The fact that the movement of a single edge is necessarily tied to a change in overall intensity would confound attempts to directly test the prediction by moving a single edge. As an alternative, we stimulated an L-neuron by movements of multiple dark and bright stripes of different spatial frequencies. If the colour of the stripes is UV, the result is as predicted (Fig. 65): the movement of the stripes causes modulation of membrane potential which differs in amplitude between upward and downward movements. For green stripes, there is no effect.

Fig. 65
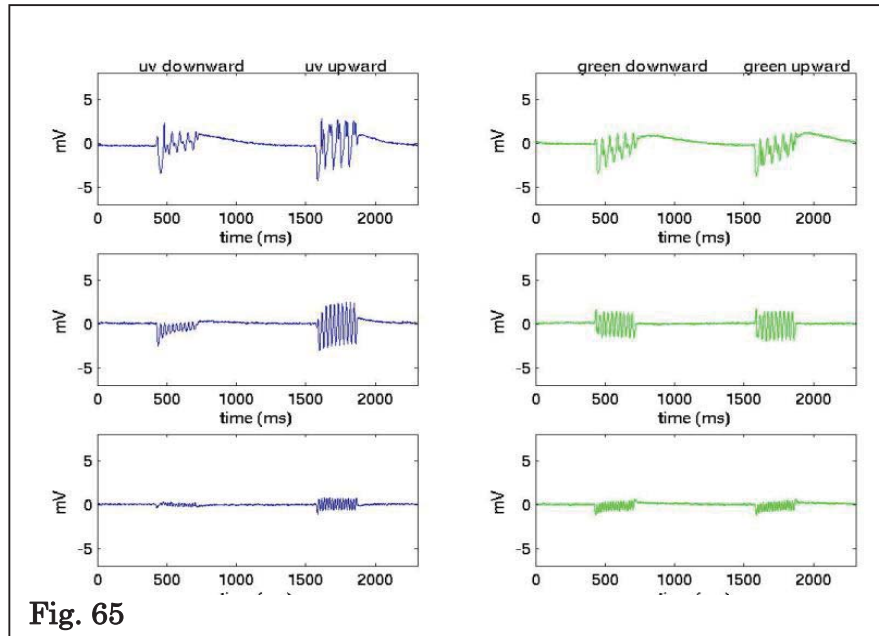
From the scientific point of view, we consider the discovery of directional sensitivity to movement as the most important single finding within the present project. We have shown that the peripheral visual system of the dragonfly ocellar system achieves a function that is akin to the function of the so-called simple cells in the human and mammalian visual cortex, but does so using far less neuronal circuitry, leading to a velocity advantage of 2 orders of magnitude. An internet link to a movie that shows movement of a cortical neuron receptive field is given in Fig. 66.


Fig. 66
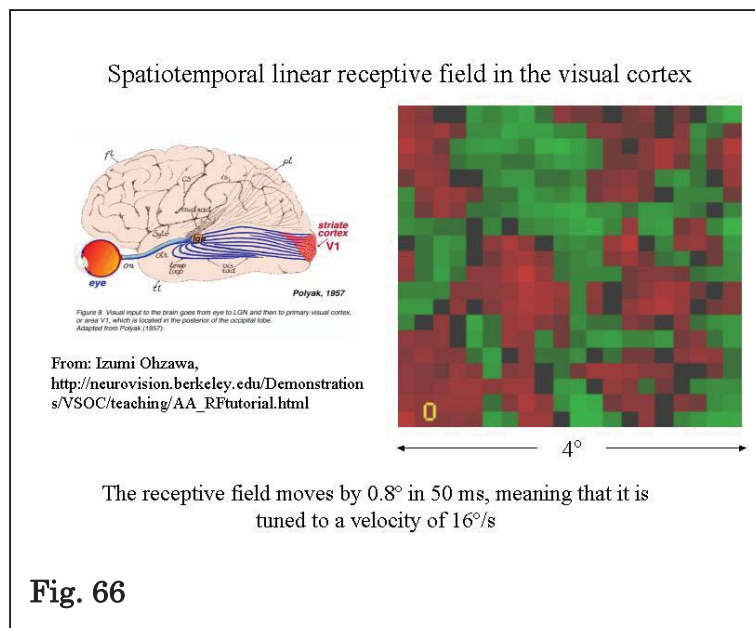
From an engineering point of view, the finding encourages us to incorporate a movement-sensitive term into an attitude control loop, in addition to the proportional and derivative terms mentioned above. None of the published attitude controllers use anything similar, leading us to the suggestion that the dragonfly has evolved a concept that is novel within the field.

33

Another look at the environment

Our investigation, in the preceding chapters, of the dragonfly ocellar system, enables us to view the environment through the filter formed by its spatiotemporal transfer characteristics. This allows us to examine the extent to which the ocellar system resoves some of the problems that are pertinent to horizon detection using any vision system:

- In an application to micro-aerial vehicles, we will sometimes be close to the ground, often encountering situations where the horizon is far from being a smooth line ('urban canyon'). A horizon detector must be reasonably immune against such effects.
- It is useful to use a wavelength range where the contrast between ground and sky is large.
- At some wavelengths, the sun forms a very high-intensity singularity. If present in a scene, which can be at a wide range of possible elevations, it may severely degrade the performance of sensors by local saturation. Therefore it is necessary to find a method whereby the effect of the sun is minimized.
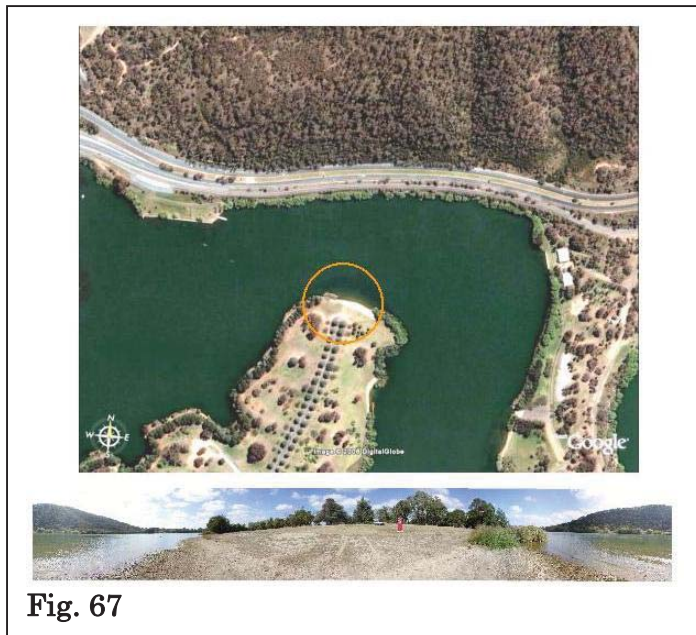
Fig. 67

The first scene (Fig 67) is a site typically inhabited by dragonflies: the shore of a peninsula. It is shown in aerial view and in a panoramic image taken from the ground, at the tip of the peninsula, 1.5 m above ground, 2 m from the shore. The panoramic image is centered on the dual row of trees on the peninsular and was, by far, not covering the full elevation range of ±90°, but rather a range of ±25°. The sun was at NNW at an elevation of 50°, which is outside the image but still sufficient to produce glare. In a first implementation of a panoramic viewer, we used a sky scanner for the measurement of UV and green photon flux at different points in any scene, consisting of single photodiodes at the ends of black tubes with apertures of 10°, mounted on a 2-D servo motor assembly. This enabled us to form a 360° by 180° panorama of the UV and green photon fluxes of the site shown above. By spatially low-pass filtering this data with a standard dragonfly ocellar photoreceptor receptive field we were able to evaluate the flux 'seen' by the photoreceptors. Fig 68 shows the filtered scans of the scene in Fig. 67.

Fig. 68
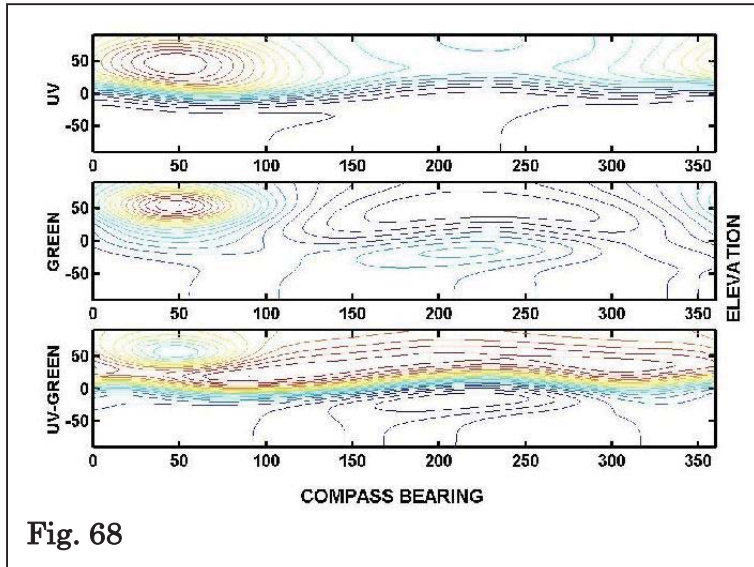
We observe that little vertical detail remains after spatial low-pass filtering with the anisotropic visual field of the ocellar photoreceptor, while the horizon is preserved. As expected, the gradient at the horizon is much more pronounced in the UV than it is in the green. In both images, the detectors saturated when the sun was in their fields of view. In the difference image the saturated patches partially cancel each other, leaving the horizon as the single most prominent feature.

Aerial view of cluttered scene (From GoogleEarth). North is up, scene is 100 m wide.

Panoramic view of scene, from location of camera, stitched from a series of pictures. The centre of the panorama faces northwest, which was also the bearing of the sun (outside field of view).

Fig. 69

Panoramic mirror

Camera

Fig. 70
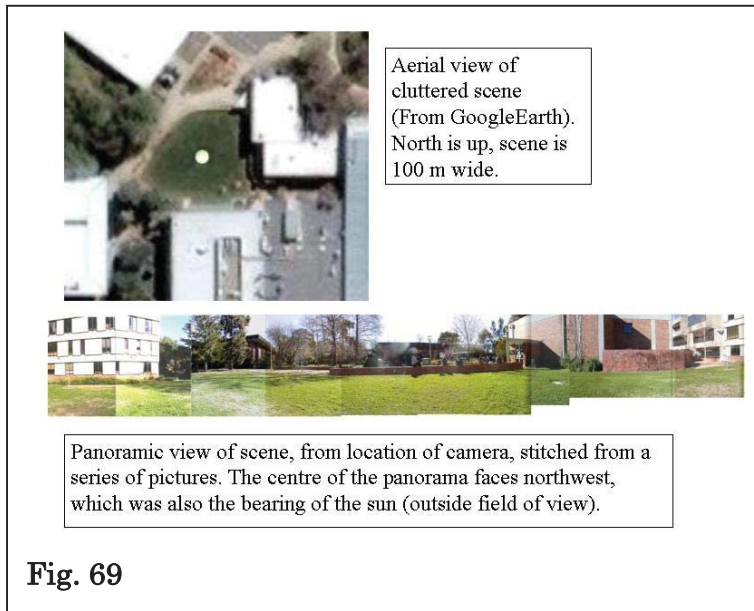
A second, more cluttered scene was a courtyard at the Australian National University (Fig. 69), surrounded by several buildings up to 3 stories high, by vegetation of various heights, and a car park on one side (up in aerial view of Fig. 69). In this case, we did not use a point scanner, but rather a UV-sensitive CCD camera, using spectral band-pass filters. As shown in Fig, 70, the camera is pointing upwards, at the surface of a parabolic mirror. The mirror is seated in a perspex ring, and its apex is pointing downwards, towards the camera. The camera was located on a bench, 1 m above ground, at the centre of the lawn-covered courtyard (Fig. 69, marked by white dot).
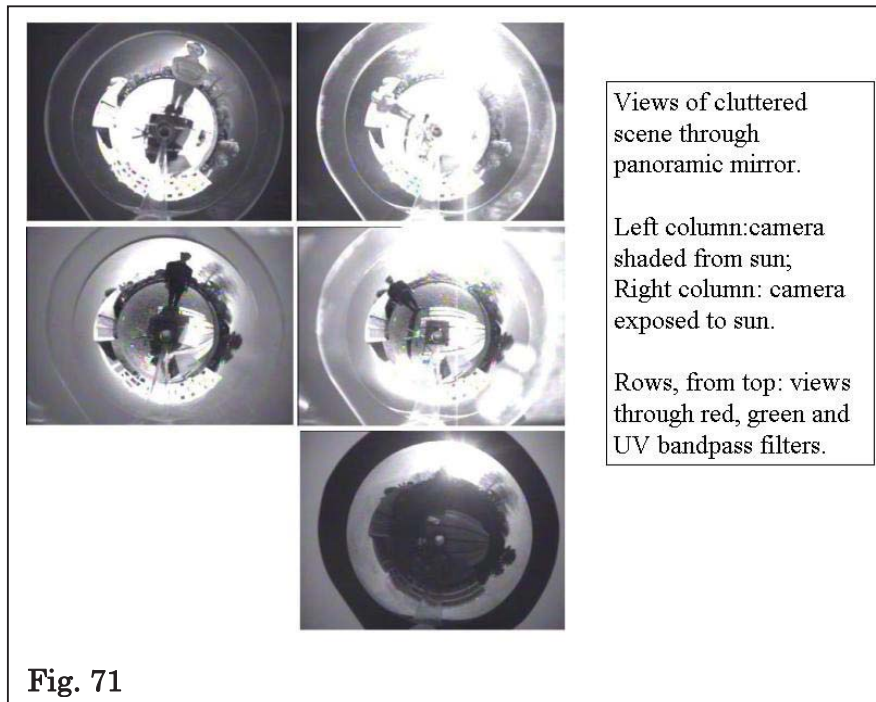
Views of cluttered scene through panoramic mirror.

Left column: camera shaded from sun; Right column: camera exposed to sun.

Rows, from top: views through red, green and UV bandpass filters.

Fig. 71

Fig 71 contains views of the scene, transformed into polar projections by the surface of the mirror, and taken at different spectral ranges. In the red (top row), the lawn and the half-ring of buildings (lower left of this image) appear bright while the half-ring of vegetation opposite the buildings appears much darker, similar to the sky. This is clearly visible only if the camera is shaded from direct sunlight (left image): with direct sun (right image), profound overexposure occurs in the sun-facing sector, rendering this image useless for horizon detection. In the green (centre row), the sky appears brighter and the ground dimmer than in the red. For the sun-exposed case, the image is even worse than in the red. In the UV (bottom), protection from the sun is not necessary to get clear contrast between sky and ground. The sky is clearly brighter than all other image elements, including buildings, trees and the lawn.

Demonstration Algorithm for Horizon Estimation

Using an image of a parabolic mirror taken by a CCD camera using a UV pass filter, the following algorithm has been used to evaluate the potential for horizon estimation using a purely visual approach in a cluttered environment.

The image to be analysed is processed in a number of steps:

- The image is broken down into a series of polar spokes radiating from the centre of the mirror. Only the section of the spoke which is between – 45° and + 45° elevation in the real world is used.

- The maximum and minimum values of the pixels along each spoke are used to normalise the values along the spoke.

- A suitable threshold value is determined and the point with the lowest elevation

which still exceeds the threshold level is noted for the next step. These low points are marked in red on the example frame. Where no threshold can be found in a particular spoke, an average radius is substituted for the following step.

- The threshold points for each spoke are low-pass filtered. This filtered threshold
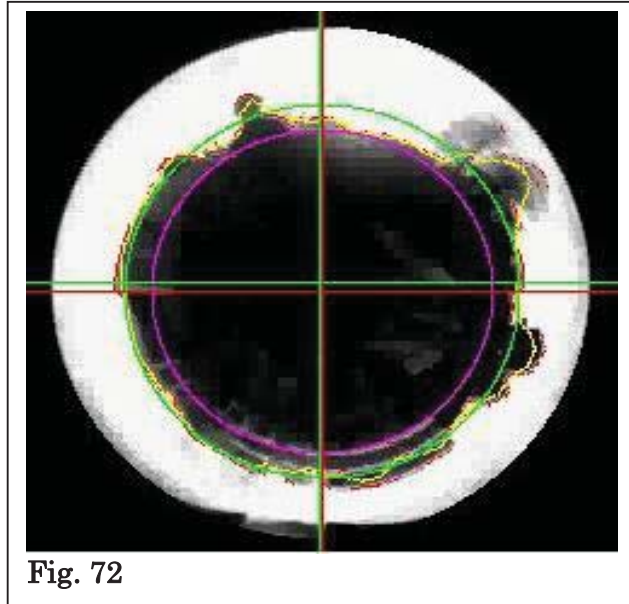
Fig. 72

line is shown in yellow on the example frame (Fig. 72).

- Now the Cartesian co-ordinates of each of these points are summed using a Fourier convolution and this gives two coordinates which represent the centre of the horizon circle. An average of the radius from the centre of the mirror to the threshold values is only used to show a green circle on the example frame. The green cross-hair on the example frame shows the mirror centre while the red cross-hair shows the estimated centre of the earth. The difference between the two cross-hairs is the calculated pitch and roll error.

- A second approximation is made by reducing the horizon radius to include the low passed threshold (the yellow line) with the lowest elevation. This is shown in purple on the example frame.

  This second approximation is still a conservative estimate in that it would tend to under-estimate a pitch and roll error.

We observe that the estimate is surprisingly accurate, inviting further validation on an extended dataset. The algorithm begins to fail where the true horizon is no longer in the field of view on all sides of the mirror.

An implementation of a machine-vision based flight attitude controller, inspired by the dragonfly ocellar system

Objectives

Within the present project we have collected a wide range of information on the structure and function of the dragonfly ocellar system. In parallel with this research, we have considered how this information might be useful for the design of an attitude controller.

This section focuses on two specific issues:

- Firstly, we propose a sensor that emulates the optical properties of the dragonfly system, allowing us to present and analyze a view of a realistic scene (as would be viewed by the dragonfly ocelli).
- Secondly, we incorporate that optical sensor into a test bed to test, in a closed loop system, spatial and temporal transfer functions and control laws as they become available from the biological research.

Design Constraints

One of the properties of the dragonfly ocellar system is its small size. This has partially motivated our study of its structure and function, with a view to an application to unmanned aerial vehicles on the micro- or even nanoscale.

The components of the ocellar system are:
- a triplet of gradient-index lenses, a few hundred microns in diameter;
- an array of approximately 2000 UV-sensitive photodetectors, each approximately 10 microns in diameter;
- a neuronal network of about twenty processing elements (second-order neurons) of a small fraction of a cubic millimeter in volume.

All those scales are within the reach of modern microelectronics and it should be possible to build a biomimetic ocellus that is not much larger than the dragonfly ocellar system, to function as the attitude controller in a dragonfly-sized helicopter. However, this would require a major effort, outside the scope of the present investigation. Consequently, we use an upscaled model that is within our reach, for the purpose of algorithm testing only.

Fig. 73

For the time being, we model the system by using a panoramic mirror, a CCD camera and a serial computer, driving an assembly of ducted fans, the latter with thrust that is sufficient to control attitude but is not intended to fully support the weight of the structure (Fig. 73). Mass and inertia of our model exceed those of the dragonfly by 3 orders of magnitude, allowing the use of a control system that is much slower. Approximate transfer functions have been modeled on measurements made from the first order (receptor) neurons and the second order (L-neurons) in the dragonfly median ocellus. The intention is to use some or all transfer functions in a real control system in order to gain insight into the ocellar system's less than obvious functions.

Our concept vehicle operates tethered to the ground by its power supply (batteries or mains) and the control system processor. Despite this we are able to experiment with individual aspects of the ocellar control system and determine an indicator as to how much high-speed control, particularly horizon detection, relies on the median and lateral ocelli. Also we may find features for noise reduction in horizon detection (e. g. sun, water reflections, clouds etc.) that is unique to a dedicated UV sensor array.

Spatial Resolution

The biological system. Taken together, the dragonfly median and lateral ocelli cover a field of view that extends over 200° in azimuth. Coverage in elevation is ±15° each above and below the horizon; large segments of the viewsphere towards the poles are not visible to the ocelli. The dragonfly's photoreceptors have resolutions of 15° in elevation and 30° in azimuth. The time resolution of the ocellar system is in the order of 10 milliseconds.

Our implementation. Using a central panoramic camera arrangement with a

parabolic convex mirror, we accomplish 360° in azimuth and 45° above and below the horizon in elevation. After unwrapping, we achieve an effective resolution of 1.4°, allowing further spatial low-pass filtering to model the low spatial resolution of the ocellar system. The spectral sensitivity can be adjusted using color filters. The image processing delay is 80 ms, an order of magnitude slower than the ocelli, but then the inertia of the model system is much larger (see above).

The basic specifications are:
Camera: Sony XC-EI30 1/3" CCD Manual exposure control.
Camera Image: 768 W x 494 H 2:1 Interlace at 25Hz (40ms/Frame)
Spectral range: UV B (>350nm) to near infrared (<900 nm)
Frame Grabber: Brooktree chipset Bt848
Frame Grabber Image: 768 W x 576 H, 2:1 Interlace at 25Hz
Software Interface image: 384 W x 288 H Alternate Fields at 50Hz
Remapped Image: 256 W x 64 H (representing 360 W x 90 H degrees)
Resolution = 1.4 degree (0.02 radian)

Temporal Resolution

Currently the sensor/processor interface operates in the following way:
in the first camera frame the CCD is exposed,
in the second frame it is transferred to the frame grabber,
in the third frame it is remapped and the control algorithm calculated
and
in the forth the control responses are sent to our control system.

Each camera field on this camera takes 20 ms so for a processing lag of 4 fields there will be a control response delay of 80 ms. The ducted fans that are used in our concept vehicle for attitude control have a similar delay due to inertia in the fan motors and in the ducted air that is moved by these fans (to be analyzed further).

By comparison, the control responses measured in our wind tunnel experiments with *Hemicordulia* and *Hemianax* dragonflies show delays (from the beginning of the visual stimulus to a measurable response) of between 50 and 100 ms. Neurological measurements indicate that the dragonfly control system could actually be a factor of 2 faster than this.

While the control system delay in both the dragonfly and the concept vehicle are comparable, their weights are not. A typical dragonfly (*Hemicordulia*) weighs about 300 milligrams and seems to have a system response of the same order as their wing beat rate (typically 20 beats per second but up to 60 beats per second). The ocellar system is mounted in a head which has some rotational freedom from the thorax and which only weighs 46 milligrams. Our concept vehicle has 4 motor/fan units each weighing 100 grams, a camera weighing 50 grams and a mirror weighing 76 grams. Total weight with the motor speed controller is about 700 grams and it has a natural angular rotation frequency of about 1 Hz. Depending of whether you use the dragonfly head or total weight, the weight ratio is between 2000 and 18000.

So, for reasons of inertia, rather than control system speed, it will be necessary to temporally scale the system by a factor of at least 100. For a factor of just 100, a 50 ms dragonfly response is scaled to a 5 second response on the concept vehicle.

Two methods of horizon detection have been investigated for initial trials with the test vehicle.

METHOD 1: SIMPLE HORIZON DETECTION.

The method relies on recognizing the horizon boundary as a dark edge when observed in the ultraviolet spectrum. The method used has been described in the above section on the demonstration algorithm for horizon estimation. This algorithm, which is computationally fast through the use of lookup tables, provides pitch and roll error.

These control errors are fed to a standard Proportional Integral Derivative control algorithm that then adjusts the ducted fans to correct the attitude of the vehicle.

Our concept vehicle is like an inverted pendulum (similar to the way the centre of lift of a helicopter is above its centre of gravity) and so had an inbuilt negative proportional response. Helicopter pilots often compare flying a helicopter to balancing a ball bearing on top of a larger ball. The real case isn't generally that bad (rotor dynamics do add some real complications) but the strong proportional feedback in the inverted pendulum creates a natural oscillation which is hard to dampen and it is easy for the novice pilot to over-control and thus increase the oscillations drastically. Experimentation with our concept vehicle quickly showed that active damping would actually improve stability while, at the same time, deteriorating the correction response time.

In this respect, our concept vehicle is very different to the dragonfly which, except in glide mode, is likely to have a more neutral or negative proportional feedback stability (more like an inverted pendulum) in order to increase agility. We can change are concept vehicle by the addition of weights to create a similar stability once some confidence is gained in the control system.

Initial trial and error selection of PID control parameters has achieved good horizon stability. Following a step response input (equivalent to a gust of wind) the platform is stabilized after 3 oscillations (about 3 seconds).

Further mathematical analysis involving the correct modelling of the concept vehicle should provide PID settings with a marginally better response rate

METHOD 2. OCELLAR HORIZON DETECTION

While the simple horizon detection method described above is good for proving the apparatus, our primary goal is to create a control system that utilises the same detection parameters as the dragonfly ocellar system. We plan to do this by creating virtual ocellar neurons in the software.

The neurons examined in the dragonfly ocelli appear to have a spatial focus that shifts temporally (or more correctly, a response delay which shifts spatially).

Predominant are two marked areas approximately 14 degrees apart that straddle the assumed position of the horizon. To simulate this in software, a second remapping of the panoramic image is made using Gaussian distributions into two virtual neurons (2 representing the two spatially separated fields in one ocellar neuron). Data collected is tabled and the measurements for the two fields at different times are summed to create a response to a horizon moving between these two visual fields.

Initial tests have shown that if this is added to the PID control loop with a positive bias, stability of the platform is achieved in about one oscillation (compared to about 3 oscillations with the PID loop).

Further Development with the Current Implementation

It is expected that insights into the stability systems in the dragonfly gained using the concept vehicle will be directly applicable to control of UAV / MAVs. As has been indicated, even initial trials have shown that a virtual ocellus can improve horizon stability and we may be very close to understanding an important flight control concept. In the near term, insights gained using the concept vehicle will be helpful in the analysis of the dragonfly control system.

Conclusions

We have started with a description of ocellar function as it was known prior to the present project, followed by a description of technical horizon detectors. A comparison leads to the recognition that there are two types of technical detectors. One is similar to the ocellar concept in using the minimum possible number of detectors, namely two, each covering a full 360º in azimuth by sinusoidal receptive fields, offset against each other by 90º. The other is based on the recognition or feature extraction of the horizon as a line, similar to the way a human pilot would identify the horizon. This approach uses a multiple pixel detector array, and in the existing implementations the field of view was quite constrained and far from panoramic.

We continued with an analysis of the dragonfly ocellar system, beginning with the optics and proceeding to the second order neurons which are the final stage of visual processing, prior to the connection to multimodal descending neurons that directly drive the dragonfly's flight motor.

We find the following salient attributes of the dragonfly system:
- The spectral sensitivity maximum is in the UV, indicating that this wavelength range is best suited for horizon detection.
- In azimuth, the total FOV is 200º. This is slightly more than a half circle and may suggest that such a FOV is optimal and that a fully panoramic FOV offers no major advantages.
- The medial 120º of this FOV is sampled by a total of 11 L-neurons, in the median ocellus, with receptive fields that partially overlap.

- The median field is supplemented by 2 L-neurons each, in the lateral ocelli, which look only at a narrow range of angles, directed sideways.
- In elevation, the FOVs of all L-neurons are quite narrow, giving the impression, in the first instance, that the system is designed to form a one-dimensional representation of the horizon.
- However, it turns out that the vertical spatial resolution, present on the level of the photoreceptors, is in fact utilized to enable the system to detect movements in elevation.

In the introduction, we had listed as potential visual cues for attitude stabilization:
1. the 'centre of gravity' of mean overall illumination,
2. the position of the boundary formed by the horizon,
3. optical flow.

We find that the dragonfly ocellar system has evolved a long way away from (1); rather, we are dealing with a combination of (2) and (3). For both types of cues, the system is remarkably parsimonious, accomplishing (2) with just 15 L-neurons, and accomplishing (3) in just a single layer of neuronal connections.

Our observations on the dragonfly encouraged us to examine a view of the world as it appears through a panoramic mirror, with a FOV that forms a ring around the equator. It is suggested that this geometry is novel with regard to horizon detection. One major advantage is that a panoramic view gives the best results if the horizon is irregular, such as in an urban canyon situation.

The next logical step will be to further test the described principles in the real world, both on our concept vehicle and in direct flight tests.

References

Chahl, J., Thakoor, S., Bouffant, N. L., Stange, G., Srinivasan, M. V., Hine, B., & Zornetzer, S. (2003). Bioinspired engineering of exploration systems: A horizon sensor/attitude reference system based on the dragonfly ocelli for Mars exploration applications. Journal of Robotic Systems, 20(1), 35-42.

Cornell,T. D.; Egan G. K. (2004) Measuring Horizon Angle from Video on a small unmanned Vehicle. 2nd International Conference on autonomous Robots and Agents. December 13-15, 2004. Palmerston North, New Zealand.

Neumann, T.R., Bulthoff, H.H., "Behavior-Oriented Vision for Biomimetic Flight Control", Proceedings of the International Workshop on Biologically-Inspired Robots, UK, August 2002, pp 196-203.

Rowell, C. H. F., Reichert, H. (1986) Three descending interneurons reporting deviation from course in the locust. II. Physiology. J Comp Physiol. A 158:775-794.

Schuppe, H.,Hengstenberg, R. (1993) Optical properties of the ocelli of *Calliphora erythrocephala* and their role in the dorsal light response. J Comp Physiol A 173:143-149.

Stange, G., Howard, J. (1979) An ocellar dorsal light response in a dragonfly. J Exp Biol 83:351-355.

Stange, G., Stowe, S., Chahl, J. S., & Massaro, T. (2002). Anisotropic imaging in the dragonfly median ocellus: A matched filter for horizon detection. Journal of Comparative Physiology A, 188, 455-467.

van Kleef J, James AC, Stange G (2005) A spatiotemporal white noise analysis of photoreceptor responses to UV and green light in the dragonfly median ocellus. J Gen Physiol 126: 481-497.

Wilson, M. (1978) The functional organization of the locust ocelli. J Comp Physiol. 124:297-316.

Winkler, S.; Schulz, H. -W.; Buschmann, M.; Kordes, T.; Vörsmann, P.: Horizon Aided Low-cost GPS/INS Integration for Autonomous Micro Air Vehicle Navigation. In: Proceedings of the 1st European Micro Air Vehicle Conference and Flight Competition (Braunschweig, Germany, July 12-13, 2004)

www.cis.udel.edu/~cer/arv/readings/paper_ettinger.pdf

www.ctie.monash.edu.au/hargrave/MECSE-4-2004.pdf

www.isr.umd.edu/Labs/CSSL/horiuchilab/pubindex.html

http://robotics.eecs.berkeley.edu/~wcwu/publications/ICRA03_Sensors.pdf.

www.futaba-rc.com/radioaccys/futm0999.html

www.fmadirect.com

www.ctie.monash.edu.au/hargrave/ horizon_sensing_autopilot.pdf